management and configuration guide

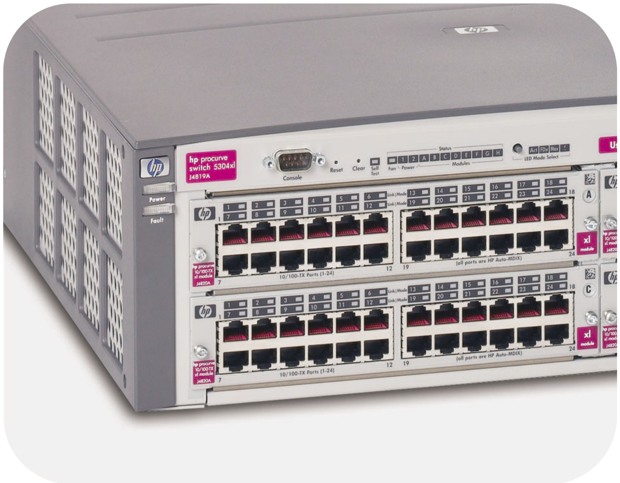

hp procurve series 5300xl switches

www.hp.com/go/hpprocurve

HP Procurve
Series 5300XL Switches

Software Release G.05.xx or Greater

Management and Configuration Guide

Publication Number

5990-3016
October 2002
Edition 4

Applicable Product

HP Procurve Switch 5308XL	(J4819A)
HP Procurve Switch 5372XL	(J4848A)
HP Procurve Switch 5348XL	(J4849A)
HP Procurve Switch 5304XL	(J4850A)

Trademark Credits

Microsoft, Windows, Windows 95, and Microsoft Windows NT are registered trademarks of Microsoft Corporation. Internet Explorer is a trademark of Microsoft Corporation. Ethernet is a registered trademark of Xerox Corporation. Netscape is a registered trademark of Netscape Corporation. Cisco® is a trademark of Cisco Systems, Inc.

Disclaimer

The information contained in this document is subject to change without notice.

HEWLETT-PACKARD COMPANY MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

Warranty

See the Customer Support/Warranty booklet included with the product.

A copy of the specific warranty terms applicable to your Hewlett-Packard products and replacement parts can be obtained from your HP Sales and Service Office or authorized dealer.

Hewlett-Packard Company
8000 Foothills Boulevard, m/s 5551
Roseville, California 95747-5551
http://www.hp.com/go/hpprocurve

Contents

Getting Started

1 Selecting a Management Interface

2 Using the Menu Interface

5 Switch Memory and Configuration

6 Interface Access, System Information, and Friendly Port Names

7 Configuring IP Addressing

8 Time Protocols

9 Optimizing Traffic Flow with Port Controls, Port Trunking, and Filters

13 802.1w Rapid Spanning Tree Protocol (RSTP) 802.1d Spanning Tree Protocol (STP)

14 Switch Meshing

16 IP Routing Features

A File Transfers

B Monitoring and Analyzing Switch Operation

C Troubleshooting

Getting Started

Contents

Overview

This *Management and Configuration Guide* is intended for use with the following switches:

- HP Procurve Switch 5304XL
- HP Procurve Switch 5308XL
- HP Procurve Switch 5348XL
- HP Procurve Switch 5372XL

Together, these four devices are termed the *Series 5300XL switches*.

This guide describes how to use the command line interface (CLI), Menu interface, and web browser interface to configure, manage, monitor, and troubleshoot switch operation. The *Product Documentation CD-ROM* shipped with the switch includes a copy of this guide. You can also download a copy from the HP Procurve website. (See "Getting Documentation From the Web" on page xxi, below.)

For information on other product documentation for Series 5300XL switch, refer to "Related Publications" on page xx.

Conventions

This guide uses the following conventions for command syntax and displayed information.

Command Syntax Statements

Syntax: aaa port-access authenticator < *port-list* >
[control < authorized | auto | unauthorized >]

- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([ ]) indicate optional elements.
- Braces (< >) enclose required elements.
- Braces within square brackets ([< >]) indicate a required element within an optional choice.
- Boldface indicates use of a CLI command, part of a CLI command syntax, or other displayed element in general text. For example:

  "Use the **copy tftp** command to download the key from a TFTP server."
- Italics indicate variables for which you must supply a value when executing the command. For example, in this command syntax, you must provide one or more port numbers:

  Syntax: aaa port-access authenticator < *port-list* >

Command Prompts

In the default configuration, your Series 5300XL switch displays one of the following CLI prompts:

```
HP Procurve Switch 5304#
HP Procurve Switch 5308#
```

To simplify recognition, this guide uses **HPswitch** to represent command prompts for all models. For example:

```
HPswitch#
```

(You can use the **hostname** command to change the text in the CLI prompt.)

Screen Simulations

Figures containing simulated screen text and command output look like this:

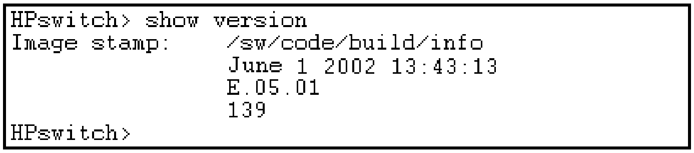

```
HPswitch> show version
Image stamp:    /sw/code/build/info
                June 1 2002 13:43:13
                E.05.01
                139
HPswitch>
```

Figure i. Example of a Figure Showing a Simulated Screen

In some cases, brief command-output sequences appear without figure identification. For example:

```
HPswitch(config)# clear public-key
HPswitch(config)# show ip client-public-key
show_client_public_key: cannot stat keyfile
```

Related Publications

Read Me First. The *Read Me First* shipped with your switch provides software update information, product notes, and other information. A printed copy is shipped with your switch. For the latest version, refer to "Getting Documentation From the Web" on page xxi.

Installation and Getting Started Guide. Use the *Installation and Getting Started Guide* shipped with your switch to prepare for and perform the physical installation. This guide also steps you through connecting the switch to your network and assigning IP addressing, as well as describing the LED indications for correct operation and trouble analysis. A PDF version of this guide is also provided on the *Product Documentation CD-ROM* shipped with the switch. And you can download a copy from the HP Procurve website. (See "Getting Documentation From the Web" on page xxi.)

Access Security Guide. Use the *Access Security Guide* to learn how to use and configure the following access security features available in the switch:

- Username and Password Security
- TACACS+ Authentication
- RADIUS Authentication and Accounting
- Secure Shell (SSH) Encryption
- Port-Based Access Control (802.1x)
- Port Security Using Authorized MAC Addresses
- Authorized IP Managers

HP provides a PDF version of this guide on the *Product Documentation CD-ROM* shipped with the switch. You can also download a copy from the HP Procurve website. (See "Getting Documentation From the Web" on page xxi.)

Command Line Interface Reference Guide. This guide, available in a PDF file on the HP Procurve website, provides a summary of the CLI commands generally available for HP Procurve switches. For the latest version, see "Getting Documentation From the Web" on page xxi.

Release Notes. Release notes are posted on the HP Procurve website and provide information on new software updates:

- New features and how to configure and use them
- Software management, including downloading software to the switch
- Software fixes addressed in current and previous releases

To view and download a copy of the latest release notes for your switch, see "Getting Documentation From the Web" on page xxi.

Getting Documentation From the Web

1. Go to the HP Procurve website at

   http://www.hp.com/go/hpprocurve

2. Click on **technical support**.
3. Click on **manuals**.
4. Click on the product for which you want to view or download a manual.

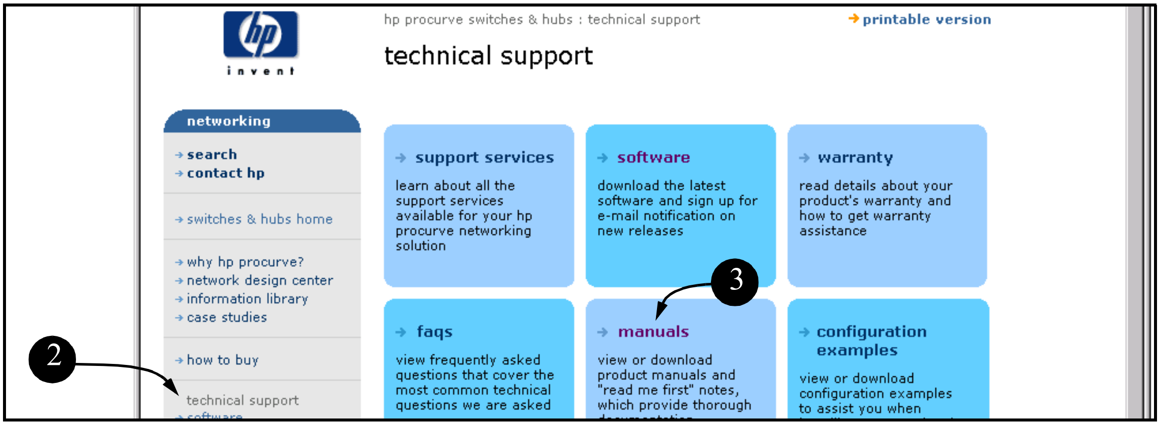

hp procurve switches & hubs : technical support

→ printable version

technical support

networking
- → search
- → contact hp
- → switches & hubs home
- → why hp procurve?
- → network design center
- → information library
- → case studies
- → how to buy
- technical support

→ **support services**
learn about all the support services available for your hp procurve networking solution

→ **software**
download the latest software and sign up for e-mail notification on new releases

→ **warranty**
read details about your product's warranty and how to get warranty assistance

→ **faqs**
view frequently asked questions that cover the most common technical questions we are asked

→ **manuals**
view or download product manuals and "read me first" notes, which provide thorough

→ **configuration examples**
view or download configuration examples to assist you when

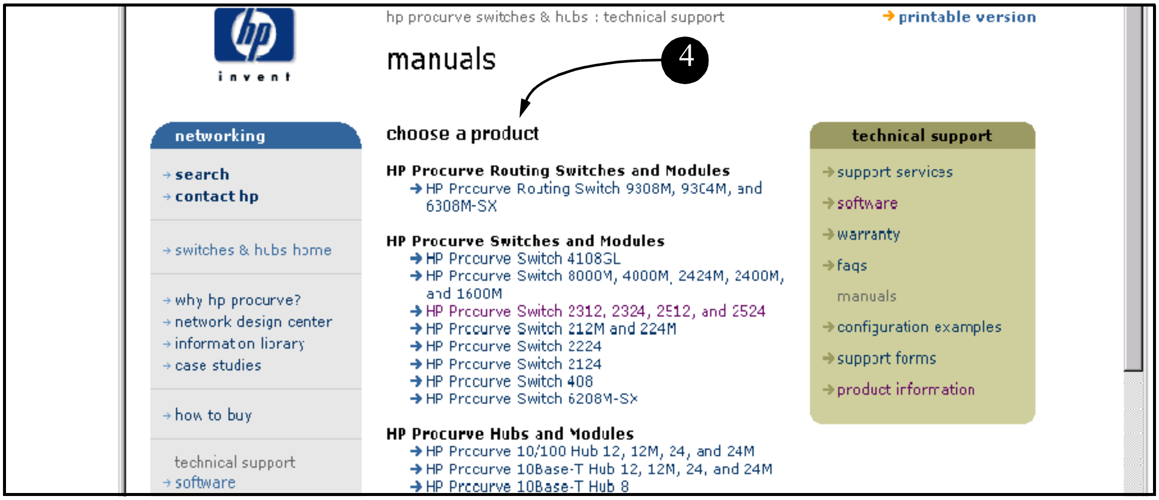

hp procurve switches & hubs : technical support

→ printable version

manuals

networking
- → search
- → contact hp
- → switches & hubs home
- → why hp procurve?
- → network design center
- → information library
- → case studies
- → how to buy
- technical support
- → software

choose a product

HP Procurve Routing Switches and Modules
- → HP Procurve Routing Switch 9308M, 9304M, and 6308M-SX

HP Procurve Switches and Modules
- → HP Procurve Switch 4108GL
- → HP Procurve Switch 8000M, 4000M, 2424M, 2400M, and 1600M
- → HP Procurve Switch 2312, 2324, 2512, and 2524
- → HP Procurve Switch 212M and 224M
- → HP Procurve Switch 2224
- → HP Procurve Switch 2124
- → HP Procurve Switch 408
- → HP Procurve Switch 6208M-SX

HP Procurve Hubs and Modules
- → HP Procurve 10/100 Hub 12, 12M, 24, and 24M
- → HP Procurve 10Base-T Hub 12, 12M, 24, and 24M
- → HP Procurve 10Base-T Hub 8

technical support
- → support services
- → software
- → warranty
- → faqs
- manuals
- → configuration examples
- → support forms
- → product information

Sources for More Information

- If you need information on specific parameters in the menu interface, refer to the online help provided in the interface. For example:

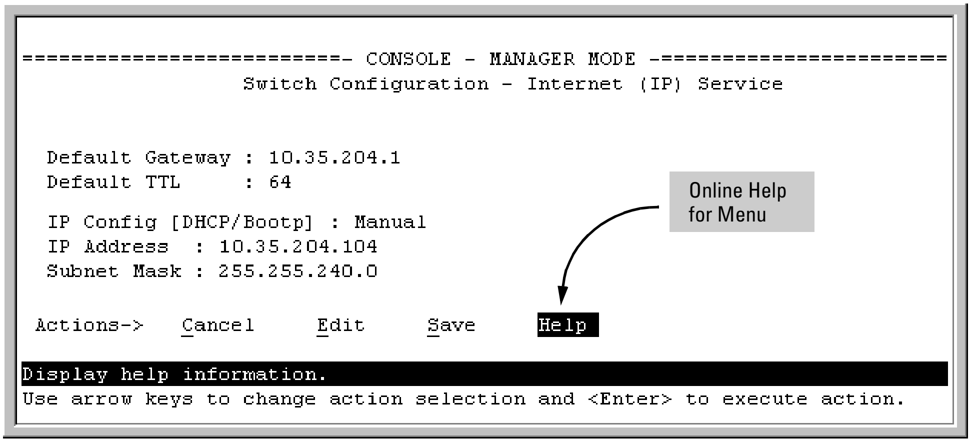

```
==========================- CONSOLE - MANAGER MODE -========================
                  Switch Configuration - Internet (IP) Service


  Default Gateway : 10.35.204.1
  Default TTL     : 64

  IP Config [DHCP/Bootp] : Manual
  IP Address  : 10.35.204.104
  Subnet Mask : 255.255.240.0

 Actions->   Cancel     Edit     Save     Help

Display help information.
Use arrow keys to change action selection and <Enter> to execute action.
```

- If you need information on a specific command in the CLI, type the command name followed by “help”. For example:

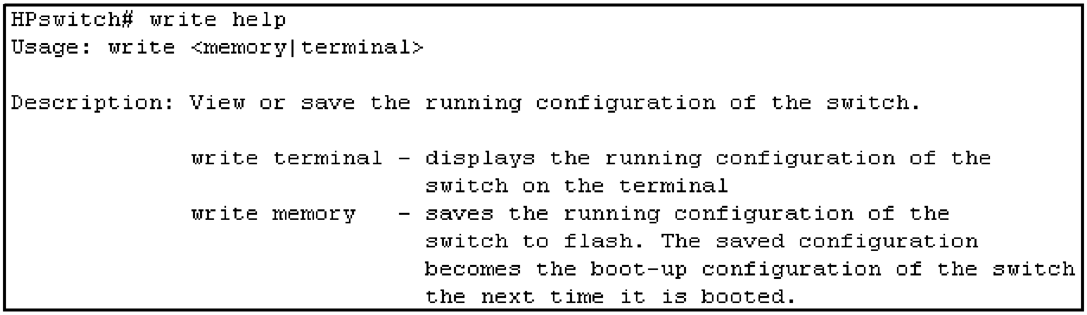

```
HPswitch# write help
Usage: write <memory|terminal>

Description: View or save the running configuration of the switch.

             write terminal - displays the running configuration of the
                              switch on the terminal
             write memory   - saves the running configuration of the
                              switch to flash. The saved configuration
                              becomes the boot-up configuration of the switch
                              the next time it is booted.
```

- If you need information on specific features in the HP Web Browser Interface (hereafter referred to as the “web browser interface”), use the online help available for the web browser interface. For more information on web browser Help options, refer to “Online Help for the HP Web Browser Interface” on page 4-11.
- If you need further information on Hewlett-Packard switch technology, visit the HP Procurve website at:

  http://www.hp.com/go/hpprocurve

Need Only a Quick Start?

IP Addressing. If you just want to give the switch an IP address so that it can communicate on your network, or if you are not using VLANs, HP recommends that you use the Switch Setup screen to quickly configure IP addressing. To do so, do one of the following:

- Enter **setup** at the CLI Manager level prompt.

```
HPswitch# setup
```

- In the Main Menu of the Menu interface, select

  8. Run Setup

For more on using the Switch Setup screen, see the *Installation and Getting Started Guide* you received with the switch.

To Set Up and Install the Switch in Your Network

Important!

Use the *HP Procurve Series 5300 Installation and Getting Started Guide* (shipped with the switch) for the following:

- Notes, cautions, and warnings related to installing and using the switch and its related modules
- Instructions for physically installing the switch in your network
- Quickly assigning an IP address and subnet mask, set a Manager password, and (optionally) configure other basic features.
- Interpreting LED behavior.

1

Selecting a Management Interface

Contents

Overview

This chapter describes the following:

- Management interfaces for the Series 5300XL switches
- Advantages of using each interface

Understanding Management Interfaces

Management interfaces enable you to reconfigure the switch and to monitor switch status and performance. The HP Series 5300XL switches offer the following interfaces:

- **Menu interface**—a menu-driven interface offering a subset of switch commands through the built-in VT-100/ANSI console—**page 1-3**
- **CLI**—a command line interface offering the full set of switch commands through the VT-100/ANSI console built into the switch—**page 1-4**
- **Web browser interface** –a switch interface offering status information and a subset of switch commands through a standard web browser (such as Netscape Navigator or Microsoft Internet Explorer)—**page 1-5**
- **HP TopTools for Hubs & Switches**–an easy-to-use, browser-based network management tool that works with HP proactive networking features built into managed HP hubs and switches

This manual describes how to use the menu interface (chapter 2), the CLI (chapter 3), the web browser interface (chapter 4), and how to use these interfaces to configure and monitor the switch.

For information on how to access the web browser interface Help, see "Online Help for the Web Browser Interface" on page 4-11.

To use HP TopTools for Hubs & Switches, refer to the *HP TopTools User's Guide* and the TopTools online help, which are available electronically with the TopTools software. (To get a copy of HP TopTools for Hubs & Switches software, see the *Read Me First* document shipped with your switch.)

Advantages of Using the Menu Interface

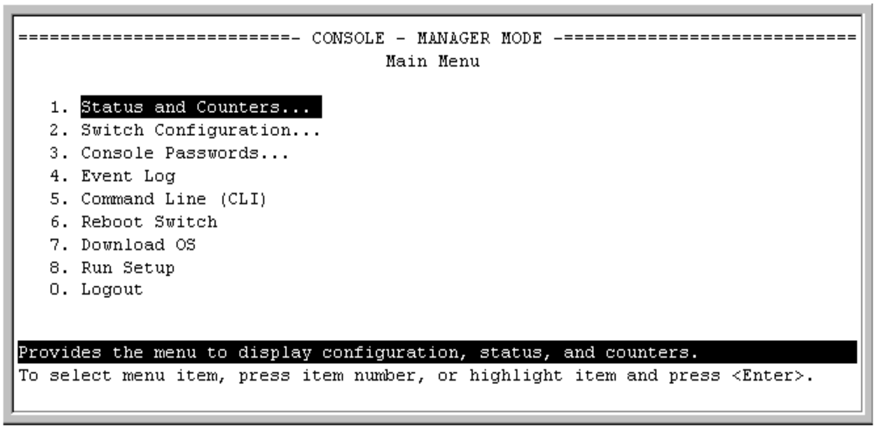

```
=========================- CONSOLE - MANAGER MODE -============================
                                  Main Menu

  1. Status and Counters...
  2. Switch Configuration...
  3. Console Passwords...
  4. Event Log
  5. Command Line (CLI)
  6. Reboot Switch
  7. Download OS
  8. Run Setup
  0. Logout


Provides the menu to display configuration, status, and counters.
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 1-1. Example of the Console Interface Display

- **Provides quick, easy management access** to a menu-driven subset of switch configuration and performance features:

  - IP addressing
  - VLANs and GVRP
  - Port Security
  - Port and Static Trunk Group
  - Spanning Tree
  - System information
  - Passwords
  - SNMP communities
  - Time protocols

  The menu interface also provides access for:

  - Setup screen
  - Event Log display
  - Switch and port status displays
  - Switch and port statistic and counter displays
  - Reboots
  - Software downloads

- **Offers out-of-band access** (through the RS-232 connection) to the switch, so network bottlenecks, crashes, lack of configured or correct IP address, and network downtime do not slow or prevent access
- **Enables Telnet (in-band) access** to the menu functionality.
- **Allows faster navigation**, avoiding delays that occur with slower display of graphical objects over a web browser interface.
- **Provides more security**; configuration information and passwords are not seen on the network.

Advantages of Using the CLI

`HPswitch>`	Operator Level
`HPswitch#`	Manager Level
`HPswitch(config)#`	Global Configuration Level
`HPswitch(`*`<context>`*`)#`	Context Configuration Levels (port, VLAN)

Figure 1-2. Command Prompt Examples

- Provides access to the complete set of the switch configuration, performance, and diagnostic features.
- Offers out-of-band access (through the RS-232 connection) or Telnet (in-band) access.
- Enables quick, detailed system configuration and management access to system operators and administrators experienced in command prompt interfaces.
- Provides help at each level for determining available options and variables.

CLI Usage

- For information on how to use the CLI, refer to chapter 3. “Using the Command Line Interface (CLI)”.
- To perform specific procedures (such as configuring IP addressing or VLANs), use the Contents listing at the front of the manual to locate the information you need.
- For monitoring and analyzing switch operation, refer to appendix B.
- For information on individual CLI commands, refer to the Index or to the online Help provided in the CLI interface.

Advantages of Using the HP Web Browser Interface

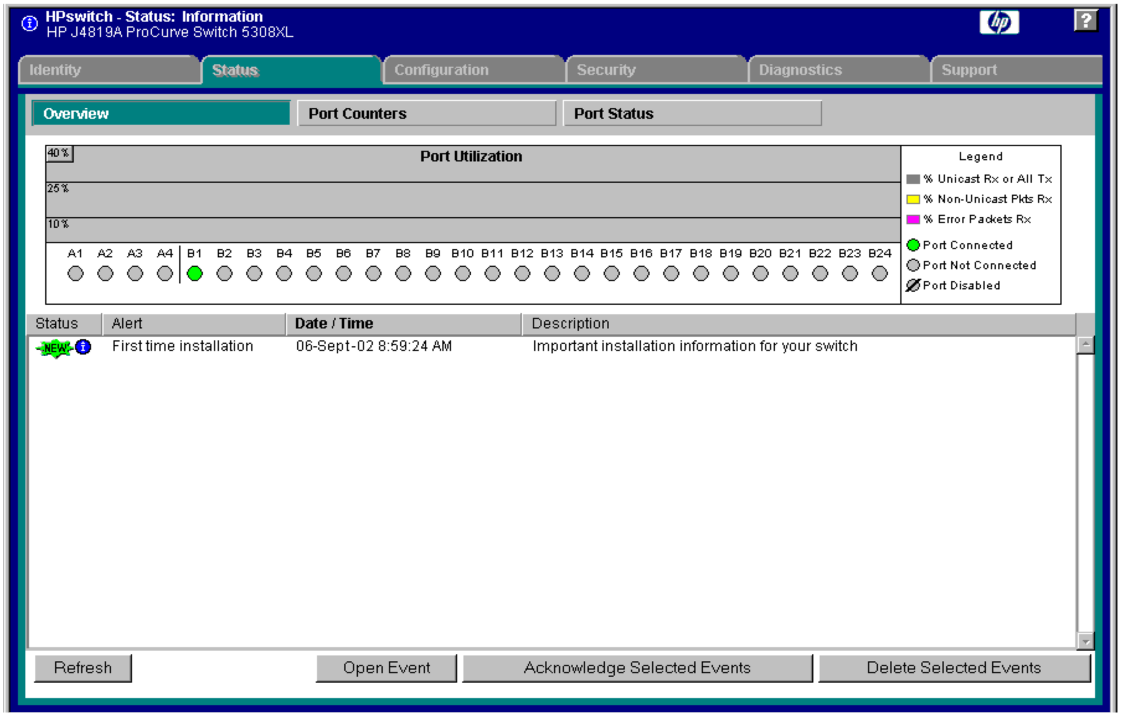

Status	Alert	Date / Time	Description
NEW i	First time installation	06-Sept-02 8:59:24 AM	Important installation information for your switch

Figure 1-3. Example of the HP Web Browser Interface

- **Easy access** to the switch from anywhere on the network
- **Familiar browser interface**–locations of window objects consistent with commonly used browsers, uses mouse clicking for navigation, no terminal setup
- **Many features have all their fields in one scree**n so you can view all values at once
- **More visual cues**, using colors, status bars, device icons, and other graphical objects instead of relying solely on alphanumeric values
- **Display of acceptable ranges of values available** in configuration list boxes

Advantages of Using HP TopTools for Hubs & Switches

You can operate HP TopTools from a PC on the network to monitor traffic, manage your hubs and switches, and proactively recommend network changes to increase network uptime and optimize performance. Easy to install and use, HP TopTools for Hubs & Switches is the answer to your management challenges.

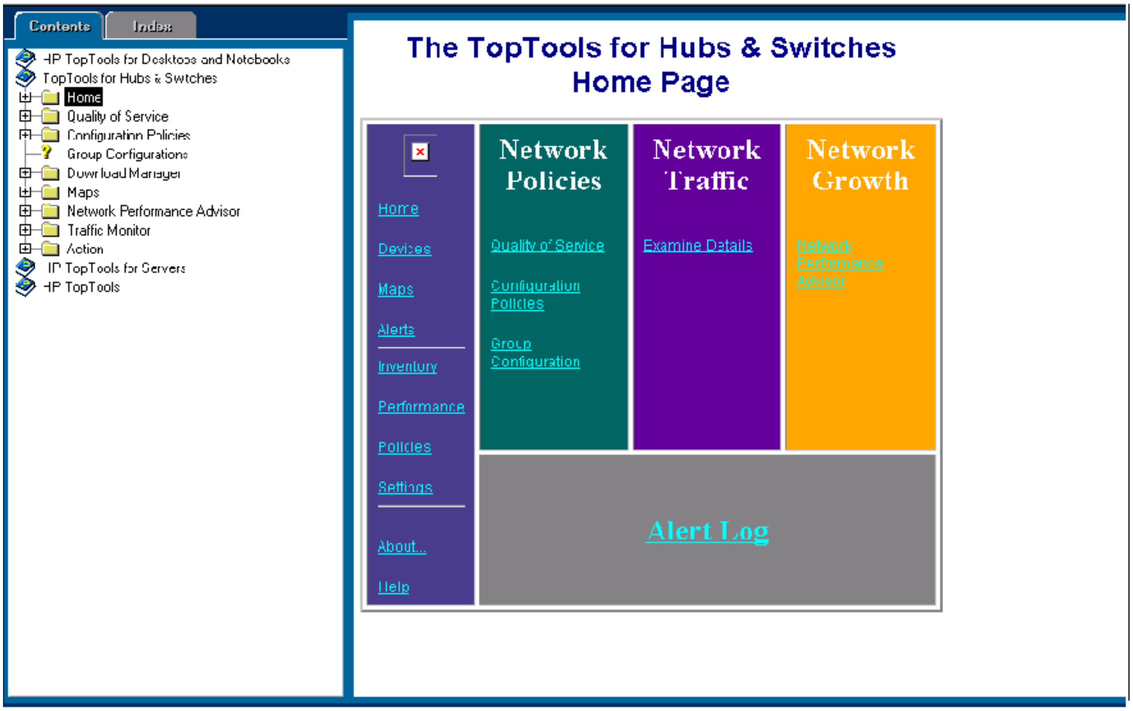

Figure 1-4. Example of HP TopTools Home Page

HP TopTools for Hubs & Switches enables greater control, uptime, and performance in your network:

- For networked devices
  - Enables fast installation of hubs and switches.
  - Enables you to proactively manage your network by using the Alert Log to quickly identify problems and suggest solutions, saving valuable time.

  - Notifies you when HP hubs use “self-healing” features to fix or limit common network problems.
  - Provides a list of discovered devices, with device type, connectivity status, the number of new or open alerts for each device, and the type of management for each device.
  - Provides graphical maps of your networked devices, from which you can access specific devices.
  - Identifies users by port and lets you assign easy-to-remember names to any network device.
  - Enables you to configure and monitor HP networked devices from your network management PC, including identity and status information, port counters, port on/off capability, sensitivity thresholds for traps, IP and security configuration, device configuration report, and other device features.
  - Enables policy-based management through the Quality of Service feature (QoS) to establish traffic priority policies for controlling and improving throughput across all the HP switches in your network that support this feature.
- For network traffic:
  - Watches the network for problems and displays real-time information about network status.
  - Shows traffic and “top talker” nodes on screen.
  - Uses traffic monitor diagrams to make bottlenecks easy to see.
  - Improves network reliability through real-time fault isolation.
  - Lets you see your entire network without having to put RMON probes on every segment (up to 1500 segments).
- For network growth:
  - Monitors, stores, and analyzes network traffic to determine where upgrades are needed.
  - Uses Network Performance Advisor for automatic traffic analysis and easy-to-understand reports that give clear, easy-to-follow plans for cost-effectively upgrading your network.

2

Using the Menu Interface

Contents

Overview

This chapter describes the following features:

- Overview of the Menu Interface (page 4-1)
- Starting and ending a Menu session (page 2-3)
- The Main Menu (page 2-7)
- Screen structure and navigation (page 2-9)
- Rebooting the switch (page 2-12)

The menu interface operates through the switch console to provide you with a subset of switch commands in an easy-to-use menu format enabling you to:

- Perform a "quick configuration" of basic parameters, such as the IP addressing needed to provide management access through your network
- Configure these features:
  - Manager and Operator passwords
  - System parameters
  - IP addressing
  - Time protocol
  - Ports
  - Trunk groups
  - A network monitoring port
  - Spanning Tree operation
  - SNMP community names
  - IP authorized managers
  - VLANs (Virtual LANs) and GVRP
- View status, counters, and Event Log information
- Update switch software
- Reboot the switch

For a detailed list of menu features, see the "Menu Features List" on page 2-14.

Privilege Levels and Password Security. HP strongly recommends that you configure a Manager password to help prevent unauthorized access to your network. A Manager password grants full read-write access to the switch. An Operator password, if configured, grants access to status and counter, Event Log, and the Operator level in the CLI. After you configure passwords on the switch and log off of the interface, access to the menu interface (and the CLI and web browser interface) will require entry of either the Manager or Operator password. (If the switch has only a Manager password, then someone without a password can still gain read-only access.)

Note

If the switch has neither a Manager nor an Operator password, anyone having access to the console interface can operate the console with full manager privileges. Also, if you configure only an Operator password, entering the Operator password enables full manager privileges.

For more information on passwords, refer to the *Access Security Guide* for your switch.

Menu Interaction with Other Interfaces.

- The menu interface displays the current running-config parameter settings. You can use the menu interface to save configuration changes made in the CLI only if the CLI changes are in the running config when you save changes made in the menu interface. (For more on how switch memory manages configuration changes, see Chapter 5, “Switch Memory and Configuration”.)
- A configuration change made through any switch interface overwrites earlier changes made through any other interface.
- The Menu Interface and the CLI (Command Line Interface) both use the switch console. To enter the menu from the CLI, use the **menu** command. To enter the CLI from the Menu interface, select **Command Line (CLI)** option.)

Starting and Ending a Menu Session

You can access the menu interface using any of the following:

- A direct serial connection to the switch’s console port, as described in the installation guide you received with the switch
- A Telnet connection to the switch console from a networked PC or the switch’s web browser interface. Telnet requires that an IP address and subnet mask compatible with your network have already been configured on the switch.

Note

This section assumes that either a terminal device is already configured and connected to the switch (see the *Installation and Getting Started Guide* shipped with your switch) or that you have already configured an IP address on the switch (required for Telnet access).

How To Start a Menu Interface Session

In its factory default configuration, the switch console starts with the CLI prompt. To use the menu interface with Manager privileges, go to the Manager level prompt and enter the **menu** command.

1. Use one of these methods to connect to the switch:
   - A PC terminal emulator or terminal
   - Telnet
2. Do one of the following:
   - If you are using Telnet, go to step 3.
   - If you are using a PC terminal emulator or a terminal, press Enter one or more times until a prompt appears.
3. When the switch screen appears, do one of the following:
   - If a password has been configured, the password prompt appears.

     ```
       Password: _
     ```

     Type the Manager password and press Enter. Entering the Manager password gives you manager-level access to the switch. (Entering the Operator password gives you operator-level access to the switch. Refer to the *Access Security Guide* for your switch.)
   - If no password has been configured, the CLI prompt appears. Go to the next step.
4. When the CLI prompt appears, display the Menu interface by entering the **menu** command. For example:

   ```
     HPswitch# menu Enter
   ```

   results in:

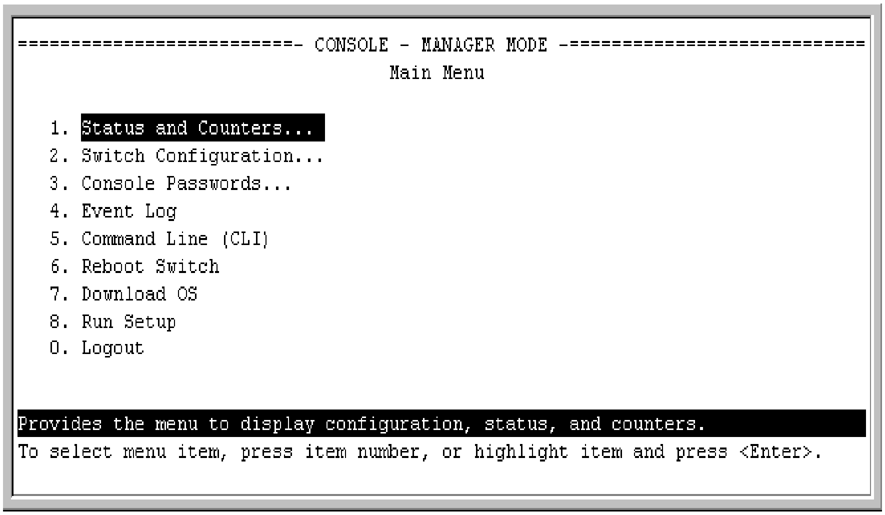

```
=========================- CONSOLE - MANAGER MODE -============================
                                  Main Menu

   1. Status and Counters...
   2. Switch Configuration...
   3. Console Passwords...
   4. Event Log
   5. Command Line (CLI)
   6. Reboot Switch
   7. Download OS
   8. Run Setup
   0. Logout


Provides the menu to display configuration, status, and counters.
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 2-1. The Main Menu with Manager Privileges

For a description of Main Menu features, see "Main Menu Features" on page 2-7.

Note

To configure the switch to start with the menu interface instead of the CLI, go to the Manager level prompt in the CLI, enter the **setup** command, and in the resulting display, change the **Logon Default** parameter to **Menu**. For more information, see the *Installation and Getting Started Guide* you received with the switch.

How To End a Menu Session and Exit from the Console:

The method for ending a menu session and exiting from the console depends on whether, during the session, you made any changes to the switch configuration that require a switch reboot to activate. (Most changes via the menu interface need only a **Save**, and do not require a switch reboot.) Configuration changes needing a reboot are marked with an asterisk (*) next to the configured item in the menu and also next to the **Switch Configuration** item in the Main Menu.

Asterisk indicates a configuration change that requires a reboot to activate.

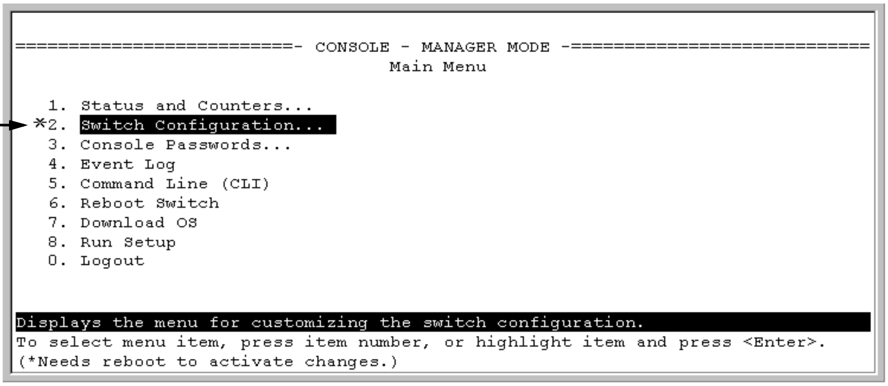

```
==========================- CONSOLE - MANAGER MODE -============================
                               Main Menu

   1. Status and Counters...
  *2. Switch Configuration...
   3. Console Passwords...
   4. Event Log
   5. Command Line (CLI)
   6. Reboot Switch
   7. Download OS
   8. Run Setup
   0. Logout


Displays the menu for customizing the switch configuration.
To select menu item, press item number, or highlight item and press <Enter>.
(*Needs reboot to activate changes.)
```

Figure 2-2. An Asterisk Indicates a Configuration Change Requiring a Reboot

1. In the current session, if you have not made configuration changes that require a switch reboot to activate, return to the Main Menu and press [0] (zero) to log out. Then just exit from the terminal program, turn off the terminal, or quit the Telnet session.

2. If you *have* made configuration changes that require a switch reboot—that is, if an asterisk (*) appears next to a configured item or next to **Switch Configuration** in the Main Menu:

   a. Return to the Main Menu.

   b. Press [6] to select **Reboot Switch** and follow the instructions on the reboot screen.

   Rebooting the switch terminates the menu session, and, if you are using Telnet, disconnects the Telnet session.

   (See “Rebooting To Activate Configuration Changes” on page 2-13.)

3. Exit from the terminal program, turn off the terminal, or close the Telnet application program.

Main Menu Features

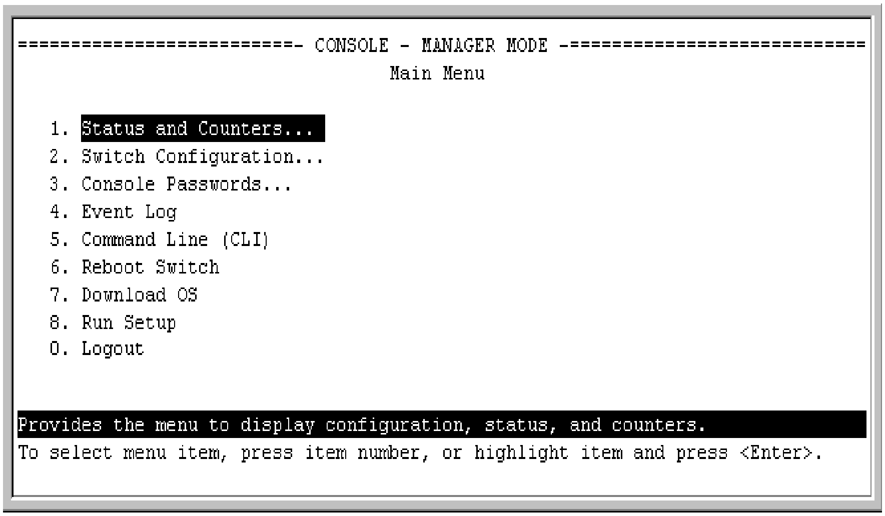

```
=========================- CONSOLE - MANAGER MODE -============================
                               Main Menu

   1. Status and Counters...
   2. Switch Configuration...
   3. Console Passwords...
   4. Event Log
   5. Command Line (CLI)
   6. Reboot Switch
   7. Download OS
   8. Run Setup
   0. Logout


Provides the menu to display configuration, status, and counters.
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 2-3. The Main Menu View with Manager Privileges

The Main Menu gives you access to these Menu interface features:

- **Status and Counters:** Provides access to display screens showing switch information, port status and counters, port and VLAN address tables, and spanning tree information. (See Chapter B, “Monitoring and Analyzing Switch Operation”.)
- **Switch Configuration:** Provides access to configuration screens for displaying and changing the current configuration settings. (See the Contents listing at the front of this manual.) For a listing of features and parameters configurable through the menu interface, see the “Menu Features List” on page 2-14 .
- **Console Passwords:** Provides access to the screen used to set or change Manager-level and Operator-level passwords, and to delete Manager and Operator password protection. (Refer to the chapter on configuring usernames and passwords in the *Access Security Guide* for your switch.)
- **Event Log:** Enables you to read progress and error messages that are useful for checking and troubleshooting switch operation. (See “Using the Event Log To Identify Problem Sources” on page C-22.)

- **Command Line (CLI):** Selects the Command Line Interface at the same level (Manager or Operator) that you are accessing in the Menu interface. (See chapter 3, "Using the Command Line Interface (CLI)".)
- **Reboot Switch:** Performs a "warm" reboot of the switch, which clears most temporary error conditions, resets the network activity counters to zero, and resets the system up-time to zero. A reboot is required to activate a change in the VLAN Support parameter. (See "Rebooting from the Menu Interface" on page 5-10.)
- **Download OS:** Enables you to download a new software version to the switch. (See Appendix A, "File Transfers".)
- **Run Setup:** Displays the Switch Setup screen for quickly configuring basic switch parameters such as IP addressing, default gateway, logon default interface, spanning tree, and others. (See the *Installation and Getting Started* guide shipped with your switch.)
- **Logout:** Closes the Menu interface and console session, and disconnects Telnet access to the switch. (See "How to End a Menu Session and Exit from the Console" on page 2-5.)

Screen Structure and Navigation

Menu interface screens include these three elements:

- Parameter fields and/or read-only information such as statistics
- Navigation and configuration actions, such as Save, Edit, and Cancel
- Help line to describe navigation options, individual parameters, and read-only data

For example, in the following System Information screen:

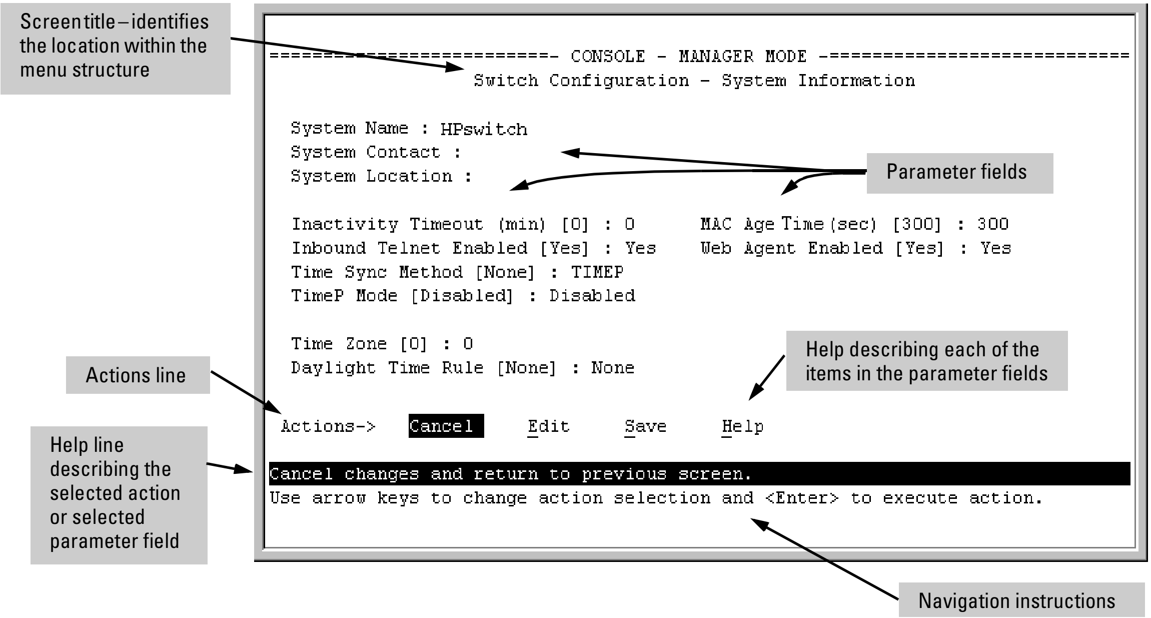

```
========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - System Information

  System Name : HPswitch
  System Contact :
  System Location :

  Inactivity Timeout (min) [0] : 0      MAC Age Time(sec) [300] : 300
  Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes
  Time Sync Method [None] : TIMEP
  TimeP Mode [Disabled] : Disabled

  Time Zone [0] : 0
  Daylight Time Rule [None] : None

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Screen title – identifies the location within the menu structure

Parameter fields

Help describing each of the items in the parameter fields

Actions line

Help line describing the selected action or selected parameter field

Navigation instructions

Figure 2-4. Elements of the Screen Structure

"Forms" Design. The configuration screens, in particular, operate similarly to a number of PC applications that use forms for data entry. When you first enter these screens, you see the current configuration for the item you have selected. To change the configuration, the basic operation is to:

1. Press [E] to select the **Edit** action.
2. Navigate through the screen making all the necessary configuration changes. (See Table 4-1 on the next page.)
3. Press [Enter] to return to the **Actions** line. From there you can save the configuration changes or cancel the changes. Cancel returns the configuration to the values you saw when you first entered the screen.

Table 2-5. How To Navigate in the Menu Interface

Task:	**Actions:**
Execute an action from the “Actions –>” list at the bottom of the screen:	Use either of the following methods: • Use the arrow keys (← ,or →) to highlight the action you want to execute, then press Enter. • Press the key corresponding to the capital letter in the action name. For example, in a configuration menu, press E to select Edit and begin editing parameter values.
Reconfigure (edit) a parameter setting or a field:	1. Select a configuration item, such as **System Name**. (See figure 2-4.) 2. Press E (for **Edit** on the Actions line). 3. Use Tab or the arrow keys (← , → , ↑ , or ↓) to highlight the item or field. 4. Do one of the following: – If the parameter has preconfigured values, either use the Space bar to select a new option or type the first part of your selection and the rest of the selection appears automatically. (The help line instructs you to “Select” a value.) – If there are no preconfigured values, type in a value (the Help line instructs you to “Enter” a value). 5. If you want to change another parameter value, return to step 3. 6. If you are finished editing parameters in the displayed screen, press Enter to return to the Actions line and do one of the following: – To save and activate configuration changes, press S (for the **Save** action). This saves the changes in the startup configuration and also implements the change in the currently running configuration. (See Chapter 5, “Switch Memory and Configuration”.) – To exit from the screen without saving any changes that you have made (or if you have not made changes), press C (for the **Cancel** action). ***Note:*** In the menu interface, executing Save activates most parameter changes and saves them in the startup configuration (or flash) memory, and it is therefore not necessary to reboot the switch after making these changes. But if an asterisk appears next to any menu item you reconfigure, the switch will not activate or save the change for that item until you reboot the switch. In this case, rebooting should be done after you have made all desired changes and then returned to the Main Menu. 7. When you finish editing parameters, return to the Main Menu. 8. If necessary, reboot the switch by highlighting Reboot Switch in the Main Menu and pressing Enter. (See the ***Note***, above.)
Exit from a read-only screen.	Press B (for the **Back** action).

To get Help on individual parameter descriptions. In most screens there is a **Help** option in the **Actions** line. Whenever any of the items in the **Actions** line is highlighted, press H, and a separate help screen is displayed. For example:

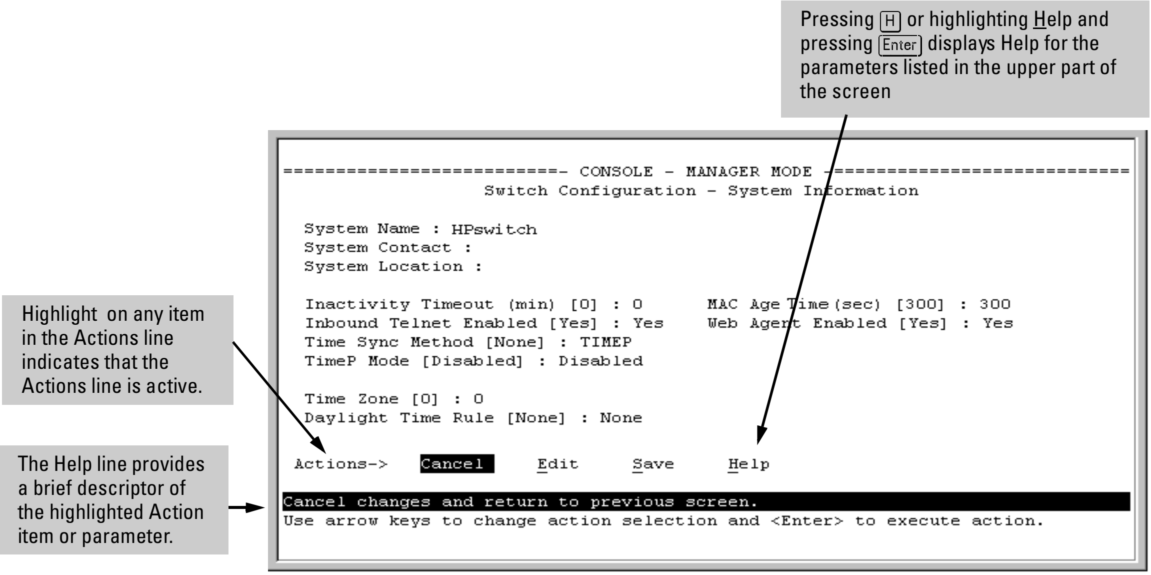

Pressing H or highlighting Help and pressing Enter displays Help for the parameters listed in the upper part of the screen

Highlight on any item in the Actions line indicates that the Actions line is active.

The Help line provides a brief descriptor of the highlighted Action item or parameter.

```
==========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - System Information

  System Name : HPswitch
  System Contact :
  System Location :

  Inactivity Timeout (min) [0] : 0       MAC Age Time(sec) [300] : 300
  Inbound Telnet Enabled [Yes] : Yes     Web Agent Enabled [Yes] : Yes
  Time Sync Method [None] : TIMEP
  TimeP Mode [Disabled] : Disabled

  Time Zone [0] : 0
  Daylight Time Rule [None] : None

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 2-6. Example Showing How To Display Help

To get Help on the actions or data fields in each screen: Use the arrow keys (←, →, ↑, or ↓) to select an action or data field. The help line under the **Actions** items describes the currently selected action or data field.

For guidance on how to navigate in a screen: See the instructions provided at the bottom of the screen, or refer to "Screen Structure and Navigation" on page 2-9.)

Rebooting the Switch

Rebooting the switch from the menu interface

- Terminates all current sessions and performs a reset of the operating system
- Activates any menu interface configuration changes that require a reboot
- Resets statistical counters to zero

(Note that statistical counters can be reset to zero without rebooting the switch.)

To Reboot the switch, use the **Reboot Switch** option in the Main Menu. (Note that the Reboot Switch option is not available if you log on in Operator mode; that is, if you enter an Operator password instead of a manager password at the password prompt.)

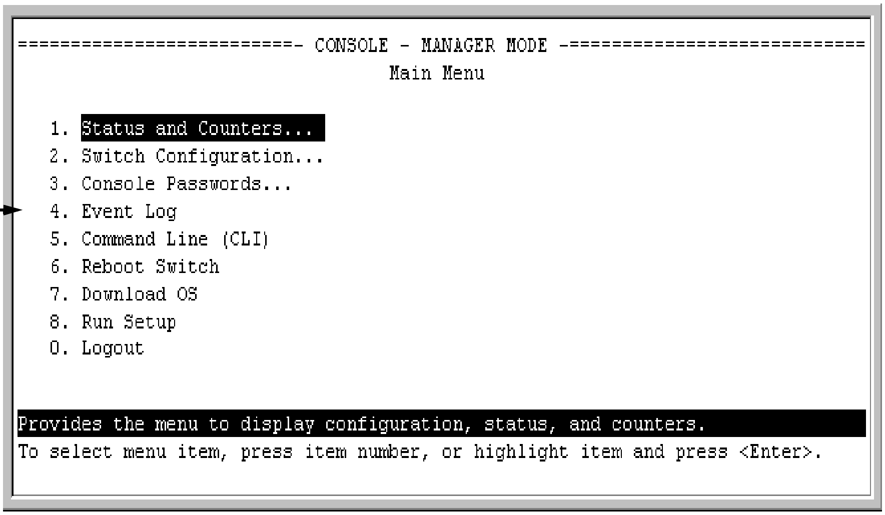

Reboot Switch option

```
==========================- CONSOLE - MANAGER MODE -============================
                                   Main Menu

   1. Status and Counters...
   2. Switch Configuration...
   3. Console Passwords...
   4. Event Log
   5. Command Line (CLI)
   6. Reboot Switch
   7. Download OS
   8. Run Setup
   0. Logout


Provides the menu to display configuration, status, and counters.
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 2-7. The Reboot Switch Option in the Main Menu

Rebooting To Activate Configuration Changes. Configuration changes for most parameters in the menu interface become effective as soon as you save them. However, you must reboot the switch in order to implement a change in the **Maximum VLANs to support parameter**. (To access this parameter, go to the Main Menu and select:

2. Switch Configuration

8. VLAN Menu

1. VLAN Support.)

If you make configuration changes in the menu interface that require a reboot, the switch displays an asterisk (*****) next to the menu item in which the change has been made. For example, if you change and save the value for the **Maximum VLANs to support** parameter, an asterisk appears next to the **VLAN Support** entry in the VLAN Menu screen, and also next to the **Switch Configuration** . . . entry in the Main Menu, as shown in figure 4-6:

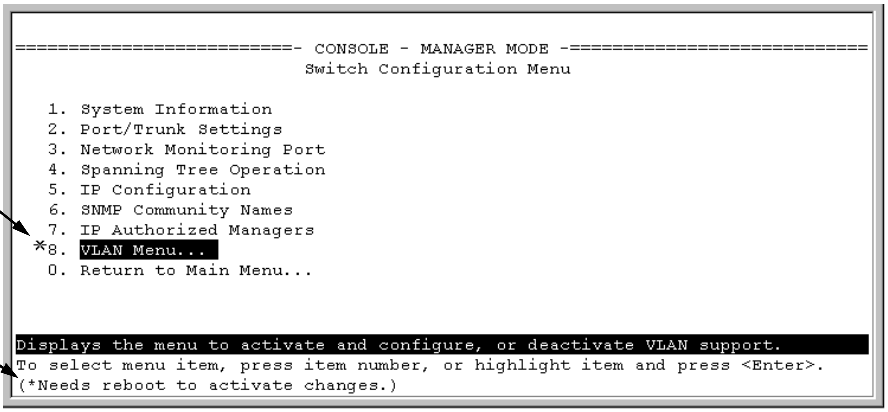

```
==========================- CONSOLE - MANAGER MODE -============================
                            Switch Configuration Menu

   1. System Information
   2. Port/Trunk Settings
   3. Network Monitoring Port
   4. Spanning Tree Operation
   5. IP Configuration
   6. SNMP Community Names
   7. IP Authorized Managers
  *8. VLAN Menu...
   0. Return to Main Menu...


Displays the menu to activate and configure, or deactivate VLAN support.
To select menu item, press item number, or highlight item and press <Enter>.
(*Needs reboot to activate changes.)
```

Asterisk indicates a configuration change that requires a reboot in order to take effect.

Reminder to reboot the switch to activate configuration changes.

Figure 2-8. Indication of a Configuration Change Requiring a Reboot

To activate changes indicated by the asterisk, go to the Main Menu and select the **Reboot Switch** option.

Note

Executing the **write memory** command in the CLI does not affect pending configuration changes indicated by an asterisk in the menu interface. That is, only a reboot from the menu interface or a **boot** or **reload** command from the CLI will activate a pending configuration change indicated by an asterisk.

Menu Features List

Status and Counters

- General System Information
- Switch Management Address Information
- Port Status
- Port Counters
- Address Table
- Port Address Table
- Spanning Tree Information

Switch Configuration

- System Information
- Port/Trunk Settings
- Network Monitoring Port
- Spanning Tree Operation
- IP Configuration
- SNMP Community Names
- IP authorized Managers
- VLAN Menu

Console Passwords

Event Log

Command Line (CLI)

Reboot Switch

Download OS

Run Setup

Logout

Where To Go From Here

This chapter provides an overview of the menu interface and how to use it. The following table indicates where to turn for detailed information on how to use the individual features available through the menu interface.

Option:	**Turn to:**
To use the Run Setup option	Refer to the *Installation and Getting Started Guide* shipped with the switch.
To view and monitor switch status and counters	Appendix B, "Monitoring and Analyzing Switch Operation"
To learn how to configure and use passwords and other security features	Refer to the *Access Security Guide* for your switch.
To learn how to use the Event Log	"Using the Event Log To Identify Problem Sources" on page C-22
To learn how the CLI operates	Chapter 3, "Using the Command Line Interface (CLI)"
To download software (the OS)	Appendix A, "File Transfers"
For a description of how switch memory handles configuration changes	"Switch Memory and Configuration" on page 5-1
For information on other switch features and how to configure them	See the Table of Contents at the front of this manual.

3

Using the Command Line Interface (CLI)

Contents

Overview

The CLI is a text-based command interface for configuring and monitoring the switch. The CLI gives you access to the switch's full set of commands while providing the same password protection that is used in the web browser interface and the menu interface.

Accessing the CLI

Like the menu interface, the CLI is accessed through the switch console, and, in the switch's factory default state, is the default interface when you start a console session. You can access the console out-of-band by directly connecting a terminal device to the switch, or in-band by using Telnet either from a terminal device or through the web browser interface.

Also, if you are using the menu interface, you can access the CLI by selecting the **Command Line (CLI)** option in the Main Menu.

Using the CLI

The CLI offers these privilege levels to help protect the switch from unauthorized access:

1. Operator
2. Manager
3. Global Configuration
4. Context Configuration

Note CLI commands are not case-sensitive.

When you use the CLI to make a configuration change, the switch writes the change to the Running-Config file in volatile memory. This allows you to test your configuration changes before making them permanent. To make changes permanent, you must use the **write memory** command to save them to the

Startup Config file in non-volatile memory. If you reboot the switch without first using **write memory**, all changes made since the last reboot or **write memory** (whichever is later) will be lost. For more on switch memory and saving configuration changes, see Chapter 5, “Switch Memory and Configuration”.

Privilege Levels at Logon

Privilege levels control the type of access to the CLI. To implement this control, you must set at least a Manager password. *Without a Manager password configured, anyone having serial port, Telnet, or web browser access to the switch can reach all CLI levels.* (For more on setting passwords, refer to the chapter on usernames and passwords in the Access Security Guide for your switch.)

When you use the CLI to log on to the switch, and passwords are set, you will be prompted to enter a password. For example:

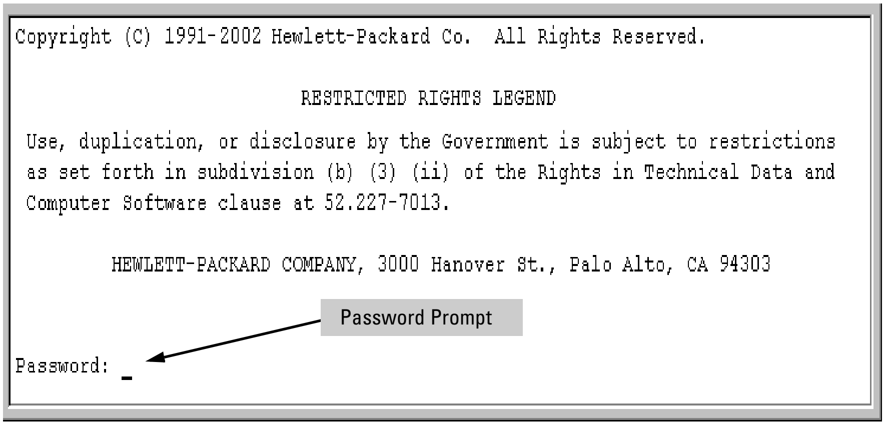

```
Copyright (C) 1991-2002 Hewlett-Packard Co.  All Rights Reserved.

                           RESTRICTED RIGHTS LEGEND

 Use, duplication, or disclosure by the Government is subject to restrictions
 as set forth in subdivision (b) (3) (ii) of the Rights in Technical Data and
 Computer Software clause at 52.227-7013.

         HEWLETT-PACKARD COMPANY, 3000 Hanover St., Palo Alto, CA 94303


Password: _
```

Figure 3-1. Example of CLI Log-On Screen with Password(s) Set

In the above case, you will enter the CLI at the level corresponding to the password you provide (operator or manager).

If no passwords are set when you log onto the CLI, you will enter at the Manager level. For example:

```
HPswitch# _
```

Caution

HP strongly recommends that you configure a Manager password. If a Manager password is not configured, then the Manager level is not password-protected, and anyone having in-band or out-of-band access to the switch may be able to reach the Manager level and compromise switch and network security. Note that configuring only an Operator password *does not* prevent access to the Manager level by intruders who have the Operator password.

Pressing the Clear button on the front of the switch removes password protection. *For this reason, it is recommended that you protect the switch from physical access by unauthorized persons*. If you are concerned about switch security and operation, you should install the switch in a secure location, such as a locked wiring closet.

Privilege Level Operation

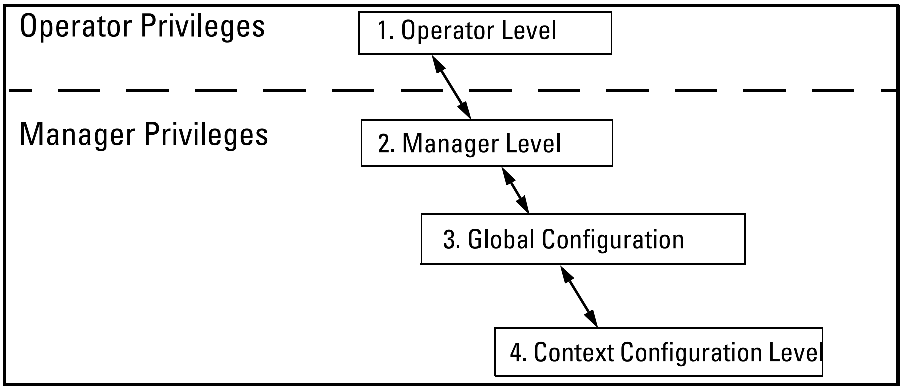

Figure 3-2. Access Sequence for Privilege Levels

Operator Privileges

At the Operator level you can examine the current configuration and move between interfaces without being able to change the configuration. A “>” character delimits the Operator-level prompt. For example:

`HPswitch> _` (Example of the Operator prompt.)

When using **enable** to move to the Manager level, the switch prompts you for the Manager password if one has already been configured.

Manager Privileges

Manager privileges give you three additional levels of access: Manager, Global Configuration, and Context Configuration. (See figure.) A "#" character delimits any Manager prompt. For example:

```
HPswitch#_             (Example of the Manager prompt.)
```

- **Manager level**: Provides all Operator level privileges plus the ability to perform system-level actions that do not require saving changes to the system configuration file. The prompt for the Manager level contains only the system name and the "#" delimiter, as shown above. To select this level, enter the **enable** command at the Operator level prompt and enter the Manager password, when prompted. For example:

  ```
  HPswitch> enable   (Enter enable at the Operator prompt.)
  HPswitch# _        (The Manager prompt.)
  ```

- **Global Configuration level:** Provides all Operator and Manager level privileges, and enables you to make configuration changes to any of the switch's software features. The prompt for the Global Configuration level includes the system name and "(config)". To select this level, enter the **config** command at the Manager prompt. For example:

  ```
  HPswitch# _        (Enter config at the Manager prompt.)
  HPswitch(config)#_(The Global Config prompt.)
  ```

- **Context Configuration level:** Provides all Operator and Manager privileges, and enables you to make configuration changes in a specific context, such as one or more ports or a VLAN. The prompt for the Context Configuration level includes the system name and the selected context. For example:

  ```
  HPswitch(eth-1)#

  HPswitch(vlan-10)#
  ```

  The Context level is useful, for example, if you want to execute several commands directed at the same port or VLAN, or if you want to shorten the command strings for a specific context area. To select this level, enter the specific context at the Global Configuration level prompt. For example, to select the context level for an existing VLAN with the VLAN ID of 10, you would enter the following command and see the indicated result:

  ```
  HPswitch(config)# vlan 10

  HPswitch(vlan-10)#
  ```

Changing Interfaces. If you change from the CLI to the menu interface, or the reverse, you will remain at the same privilege level. For example, entering the `menu` command from the Operator level of the CLI takes you to the Operator privilege level in the menu interface.

Table 3-1. Privilege Level Hierarchy

Privilege Level	Example of Prompt and Permitted Operations		
Operator Privilege			
Operator Level	`HPswitch>`	show <command> setup	View status and configuration information.
		ping <argument> link-test <argument>	Perform connectivity tests.
		enable	Move from the Operator level to the Manager level.
		menu	Move from the CLI interface to the menu interface.
		logout	Exit from the CLI interface and terminate the console session.
		exit	Terminate the current session (same as logout).
Manager Privilege			
Manager Level	`HPswitch#`	Perform system-level actions such as system control, monitoring, and diagnostic commands, plus any of the Operator-level commands. For a list of available commands, enter **?** at the prompt.	
Global Configuration Level	`HPswitch(config)#`	Execute configuration commands, plus all Operator and Manager commands. For a list of available commands, enter **?** at the prompt.	

Context Configuration Level	`HPswitch(eth-5)#`
	`HPswitch(vlan-100)#`
		Execute context-specific configuration commands, such as a particular VLAN or switch port. This is useful for shortening the command strings you type, and for entering a series of commands for the same context. For a list of available commands, enter **?** at the prompt.	

How To Move Between Levels

Change in Levels	Example of Prompt, Command, and Result

Operator level
to
Manager level
	`HPswitch> enable`
	`Password:_`
	After you enter **enable**, the Password prompt appears. After you enter the Manager password, the system prompt appears with the **#** symbol:
	`HPswitch#_`

Manager level
to
Global configuration level
	`HPswitch# config`
	`HPswitch(config)#`

Global configuration level
to a
Context configuration level
	`HPswitch(config)# vlan 10`
	`HPswitch(vlan-10)#`

Context configuration level
to another
Context configuration level
	`HPswitch(vlan-10)# interface e 3`
	`HPswitch(int-3)#`
	The CLI accepts "`e`" as the abbreviated form of "`ethernet`".

Move from any level to the preceding level	`HPswitch(int-3)# exit`
	`HPswitch(config)# exit`
	`HPswitch# exit`
	`HPswitch>`

Move from any level to the Manager level	`HPswitch(int-3)# end`
	`HPswitch#`
	—or—
	`HPswitch(config)# end`
	`HPswitch#`

Moving Between the CLI and the Menu Interface. When moving between interfaces, the switch retains the current privilege level (Manager or Operator). That is, if you are at the Operator level in the menu and select the **Command Line Interface (CLI)** option from the Main Menu, the CLI prompt appears at the Operator level.

Changing Parameter Settings. Regardless of which interface is used (CLI, menu interface, or web browser interface), the most recently configured version of a parameter setting overrides any earlier settings for that parameter. For example, if you use the menu interface to configure an IP address of "*X*" for VLAN 1 and later use the CLI to configure a different IP address of "*Y*" for VLAN 1, then "*Y*" replaces "*X*" as the IP address for VLAN 1 in the running-config file. If you subsequently execute **write memory** in the CLI, then the switch also stores "*Y*" as the IP address for VLAN 1 in the startup-config file. (For more on the startup-config and running config files, see Chapter 5, "Switch Memory and Configuration".)

Listing Commands and Command Options

At any privilege level you can:

- List all of the commands available at that level
- List the options for a specific command

Listing Commands Available at Any Privilege Level

At a given privilege level you can list and execute the commands that level offers, plus all of the commands available at preceding levels. For example, at the Operator level, you can list and execute only the Operator level commands. However, at the Manager level, you can list and execute the commands available at both the Operator and Manager levels.

Type "?" To List Available Commands. 1.Typing the **?** symbol lists the commands you can execute at the current privilege level. For example, typing **?** at the Operator level produces this listing:

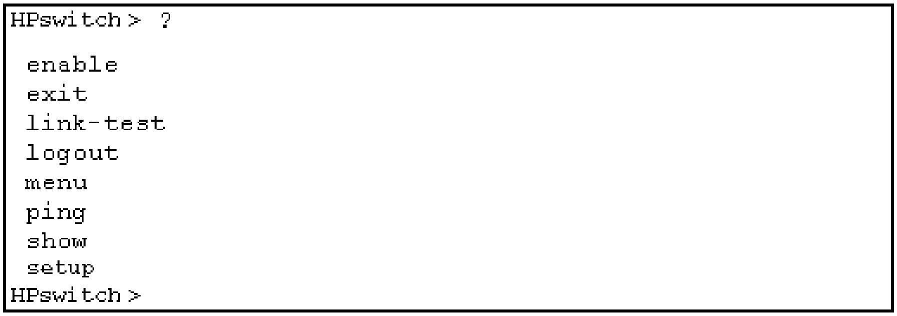

```
HPswitch> ?

 enable
 exit
 link-test
 logout
 menu
 ping
 show
 setup
HPswitch>
```

Figure 3-3. Example of the Operator Level Command Listing

Typing **?** at the Manager level produces this listing:

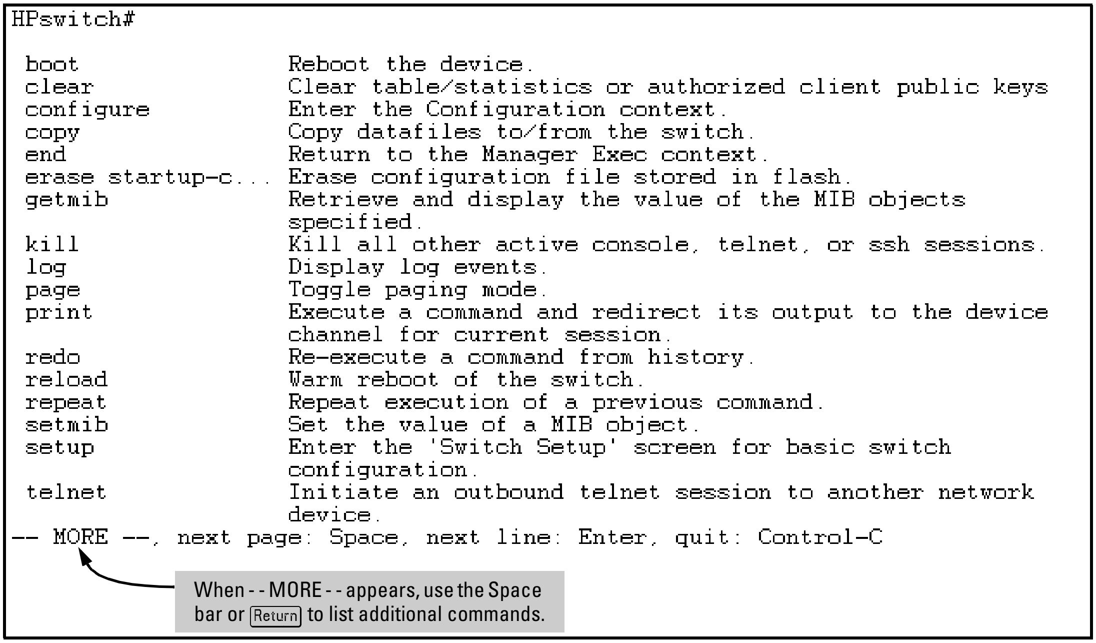

```
HPswitch#

 boot                 Reboot the device.
 clear                Clear table/statistics or authorized client public keys
 configure            Enter the Configuration context.
 copy                 Copy datafiles to/from the switch.
 end                  Return to the Manager Exec context.
 erase startup-c...   Erase configuration file stored in flash.
 getmib               Retrieve and display the value of the MIB objects
                      specified.
 kill                 Kill all other active console, telnet, or ssh sessions.
 log                  Display log events.
 page                 Toggle paging mode.
 print                Execute a command and redirect its output to the device
                      channel for current session.
 redo                 Re-execute a command from history.
 reload               Warm reboot of the switch.
 repeat               Repeat execution of a previous command.
 setmib               Set the value of a MIB object.
 setup                Enter the 'Switch Setup' screen for basic switch
                      configuration.
 telnet               Initiate an outbound telnet session to another network
                      device.
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

When - - MORE - - appears, use the Space bar or [Return] to list additional commands.

Figure 3-4. Example of the Manager-Level Command Listing

When **- - MORE - -** appears, there are more commands in the listing. To list the next screenfull of commands, press the Space bar. To list the remaining commands one-by-one, repeatedly press [Enter].

Typing **?** at the Global Configuration level or the Context Configuration level produces similar results.

Use [Tab] To Search for or Complete a Command Word. You can use [Tab] to help you find CLI commands or to quickly complete the current word in a command. To do so, type one or more consecutive characters in a command and then press [Tab] (with no spaces allowed). For example, at the Global Configuration level, if you press [Tab] immediately after typing "**t**", the CLI displays the available command options that begin with "t". For example:

```
HPswitch(config)# t[Tab]
telnet-server
time
trunk
telnet
terminal
HPswitch(config)# t
```

As mentioned above, if you type part of a command word and press [Tab], the CLI completes the current word (if you have typed enough of the word for the CLI to distinguish it from other possibilities), including hyphenated extensions. For example:

```
HPswitch(config)# port[Tab]
HPswitch(config)# port-security _
```

Pressing [Tab] after a completed command word lists the further options for that command.

```
HPswitch(config)# qos[Tab]

 device-priority      Configure device-based priority.
 dscp-map             Define mapping between a DSCP
                      (Differentiated-Services
                      Codepoint) value and 802.1p
                      priority.
 protocol             Configure protocol-based
                      priority.
 udp-port             Set UDP port based priority.
 tcp-port             Set TCP port based priority.

 type-of-service      Configure the Type-of-Service
                      method the device uses to
                      prioritize IP traffic.
```

Listing Command Options

You can use the CLI to remind you of the options available for a command by entering command keywords followed by **?**. For example, suppose you want to see the command options for configuring port C5:

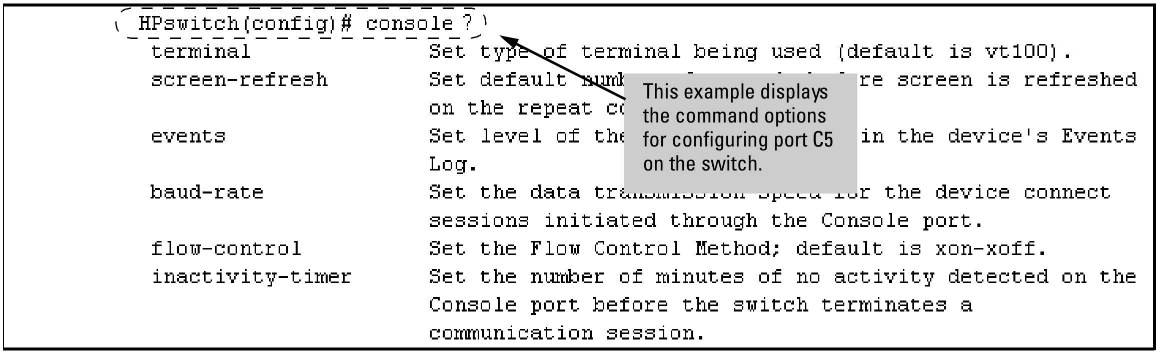

```
HPswitch(config)# console ?
  terminal              Set type of terminal being used (default is vt100).
  screen-refresh        Set default numb                re screen is refreshed
                        on the repeat co
  events                Set level of the                in the device's Events
                        Log.
  baud-rate             Set the data transmission speed for the device connect
                        sessions initiated through the Console port.
  flow-control          Set the Flow Control Method; default is xon-xoff.
  inactivity-timer      Set the number of minutes of no activity detected on the
                        Console port before the switch terminates a
                        communication session.
```

This example displays the command options for configuring port C5 on the switch.

Figure 3-5. Example of How To List the Options for a Specific Command

Displaying CLI “Help”

CLI Help provides two types of context-sensitive information:

- Command list with a brief summary of each command’s purpose
- Detailed information on how to use individual commands

Displaying Command-List Help. You can display a listing of command Help summaries for all commands available at the current privilege level. That is, when you are at the Operator level, you can display the Help summaries only for Operator-Level commands. At the Manager level, you can display the Help summaries for both the Operator and Manager levels, and so on.

Syntax: help

For example, to list the Operator-Level commands with their purposes:

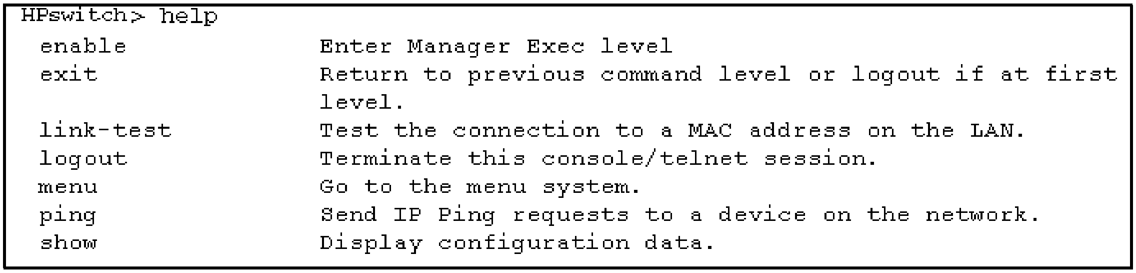

```
HPswitch> help
 enable                  Enter Manager Exec level
 exit                    Return to previous command level or logout if at first
                         level.
 link-test               Test the connection to a MAC address on the LAN.
 logout                  Terminate this console/telnet session.
 menu                    Go to the menu system.
 ping                    Send IP Ping requests to a device on the network.
 show                    Display configuration data.
```

Figure 3-6. Example of Context-Sensitive Command-List Help

Displaying Help for an Individual Command. You can display Help for any command that is available at the current context level by entering enough of the command string to identify the command, along with help.

Syntax: *<command string>* help

For example, to list the Help for the **interface** command in the Global Configuration privilege level:

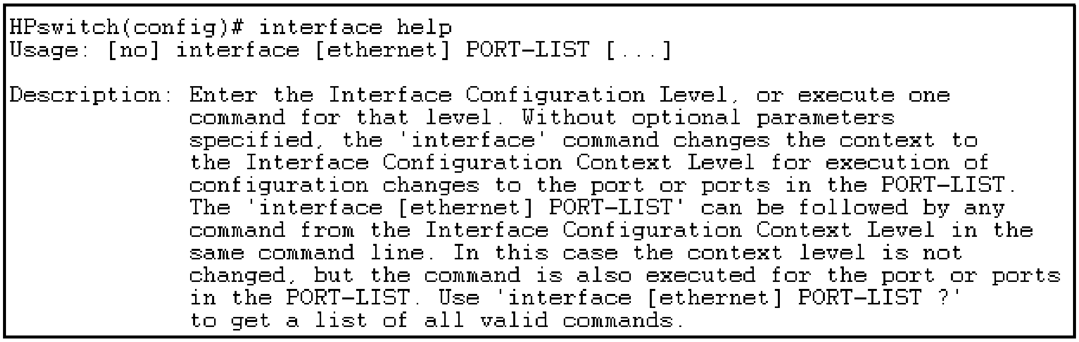

```
HPswitch(config)# interface help
Usage: [no] interface [ethernet] PORT-LIST [...]

Description: Enter the Interface Configuration Level, or execute one
             command for that level. Without optional parameters
             specified, the 'interface' command changes the context to
             the Interface Configuration Context Level for execution of
             configuration changes to the port or ports in the PORT-LIST.
             The 'interface [ethernet] PORT-LIST' can be followed by any
             command from the Interface Configuration Context Level in the
             same command line. In this case the context level is not
             changed, but the command is also executed for the port or ports
             in the PORT-LIST. Use 'interface [ethernet] PORT-LIST ?'
             to get a list of all valid commands.
```

Figure 3-7. Example of How To Display Help for a Specific Command

A similar action lists the Help showing additional parameter options for a given command. The following example illustrates how to list the Help for an interface command acting on a specific port:

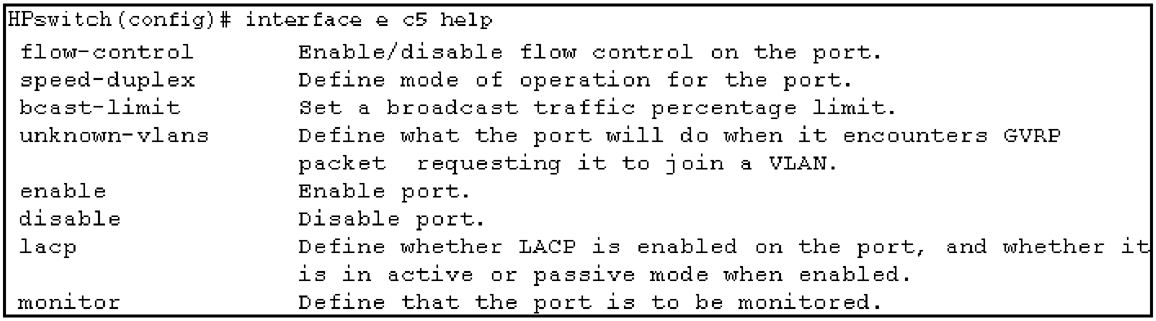

```
HPswitch(config)# interface e c5 help
 flow-control        Enable/disable flow control on the port.
 speed-duplex        Define mode of operation for the port.
 bcast-limit         Set a broadcast traffic percentage limit.
 unknown-vlans       Define what the port will do when it encounters GVRP
                     packet  requesting it to join a VLAN.
 enable              Enable port.
 disable             Disable port.
 lacp                Define whether LACP is enabled on the port, and whether it
                     is in active or passive mode when enabled.
 monitor             Define that the port is to be monitored.
```

Figure 3-8. Example of Help for a Specific Instance of a Command

Note that trying to list the help for an individual command from a privilege level that does not include that command results in an error message. For example, trying to list the help for the **interface** command while at the global configuration level produces this result:

```
HPswitch# interface help
Invalid input: interface
```

Configuration Commands and the Context Configuration Modes

You can execute any configuration command in the global configuration mode or in selected context modes. However, using a context mode enables you to execute context-specific commands faster, with shorter command strings.

The Series 5300XL switches offers interface (port or trunk group) and VLAN context configuration modes:

Port or Trunk-Group Context . Includes port- or trunk-specific commands that apply only to the selected port(s) or trunk group, plus the global configuration, Manager, and Operator commands. The prompt for this mode includes the identity of the selected port(s):

`HPswitch(config)# interface e c3-c6`
`HPswitch(config)# interface e trk1`
	Command executed at configuration level for entering port or trk1 static trunk-group context.

`HPswitch(eth-C5-C8)#`
`HPswitch(eth-Trk1)#`
	Resulting prompt showing port or static trunk contexts.

`HPswitch(eth-C5-C8)# ?` `HPswitch(eth-C5-C8)# ?`	Lists the commands you can use in the port or static trunk context, plus the Manager, Operator, and context commands you can execute at this level.

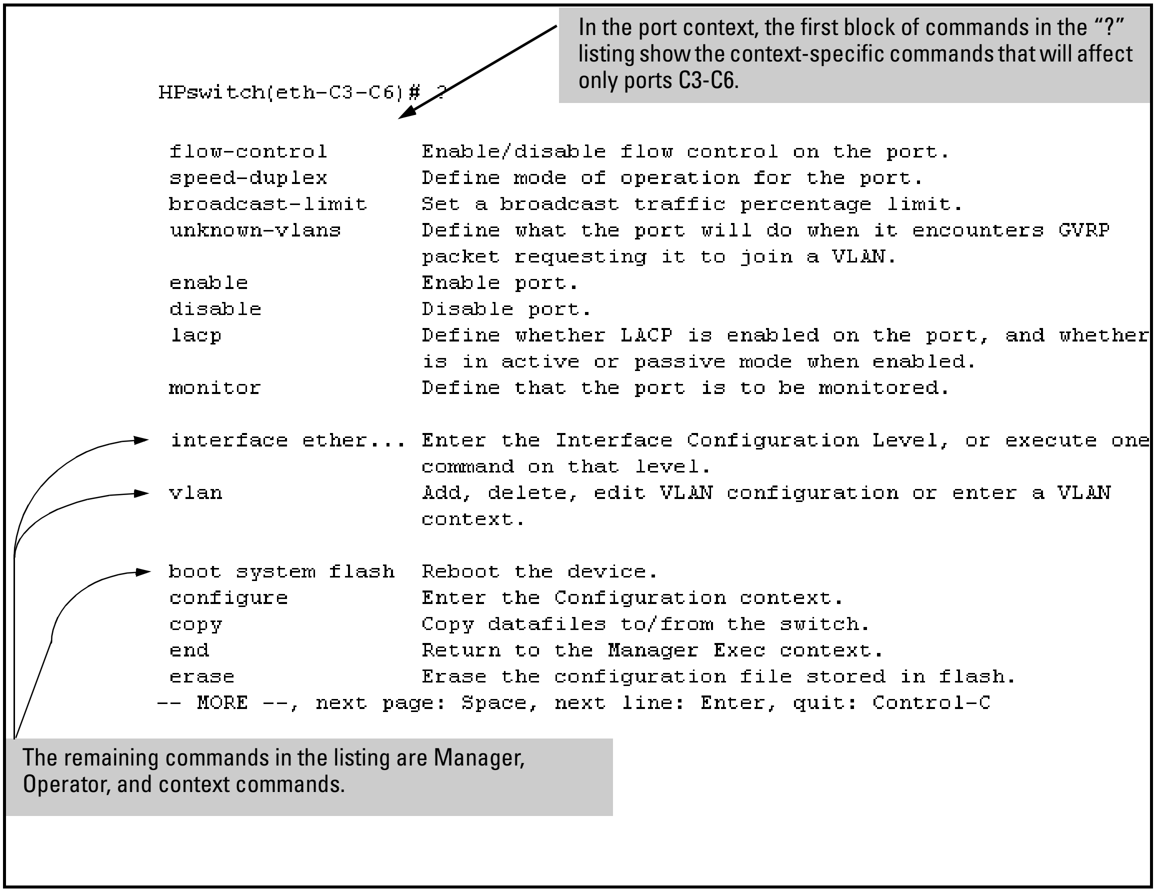

```
HPswitch(eth-C3-C6)# ?

 flow-control       Enable/disable flow control on the port.
 speed-duplex       Define mode of operation for the port.
 broadcast-limit    Set a broadcast traffic percentage limit.
 unknown-vlans      Define what the port will do when it encounters GVRP
                    packet requesting it to join a VLAN.
 enable             Enable port.
 disable            Disable port.
 lacp               Define whether LACP is enabled on the port, and whether
                    is in active or passive mode when enabled.
 monitor            Define that the port is to be monitored.

 interface ether... Enter the Interface Configuration Level, or execute one
                    command on that level.
 vlan               Add, delete, edit VLAN configuration or enter a VLAN
                    context.

 boot system flash  Reboot the device.
 configure          Enter the Configuration context.
 copy               Copy datafiles to/from the switch.
 end                Return to the Manager Exec context.
 erase              Erase the configuration file stored in flash.
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

In the port context, the first block of commands in the “?” listing show the context-specific commands that will affect only ports C3-C6.

The remaining commands in the listing are Manager, Operator, and context commands.

Figure 3-9. Context-Specific Commands Affecting Port Context

VLAN Context . Includes VLAN-specific commands that apply only to the selected VLAN, plus Manager and Operator commands. The prompt for this mode includes the VLAN ID of the selected VLAN. For example, if you had already configured a VLAN with an ID of 100 in the switch:

`HPswitch(config)# vlan 100`	Command executed at configuration level to enter VLAN 100 context.
`HPswitch(vlan-100)#`	Resulting prompt showing VLAN 100 context.
`HPswitch(vlan-100)# ?`	Lists commands you can use in the VLAN context, plus Manager, Operator, and context commands you can execute at this level.

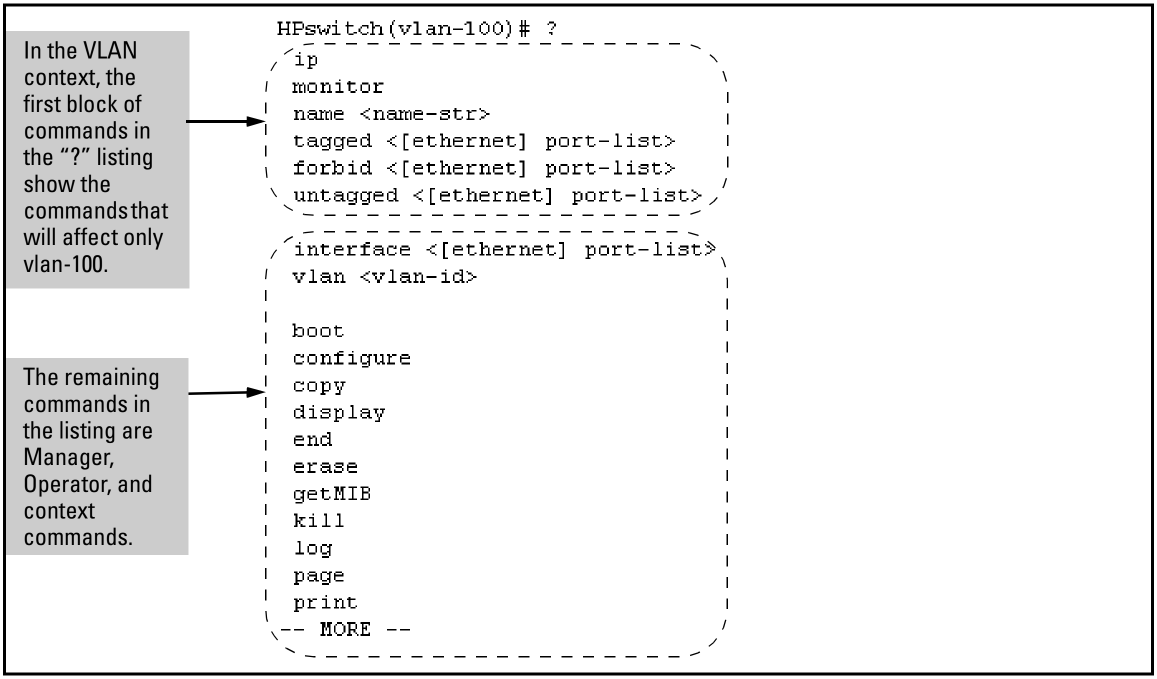

```
HPswitch(vlan-100)# ?
  ip
  monitor
  name <name-str>
  tagged <[ethernet] port-list>
  forbid <[ethernet] port-list>
  untagged <[ethernet] port-list>

  interface <[ethernet] port-list>
  vlan <vlan-id>

  boot
  configure
  copy
  display
  end
  erase
  getMIB
  kill
  log
  page
  print
-- MORE --
```

In the VLAN context, the first block of commands in the "?" listing show the commands that will affect only vlan-100.

The remaining commands in the listing are Manager, Operator, and context commands.

Figure 3-10. Context-Specific Commands Affecting VLAN Context

CLI Control and Editing

Keystrokes	Function
Ctrl A	Jumps to the first character of the command line.
Ctrl B or ←	Moves the cursor back one character.
Ctrl C	Terminates a task and displays the command prompt.
Ctrl D	Deletes the character at the cursor.
Ctrl E	Jumps to the end of the current command line.
Ctrl F or →	Moves the cursor forward one character.
Ctrl K	Deletes from the cursor to the end of the command line.
Ctrl L or Ctrl R	Repeats current command line on a new line.
Ctrl N or ↓	Enters the next command line in the history buffer.
Ctrl P or ↑	Enters the previous command line in the history buffer.
Ctrl U or Ctrl X	Deletes from the cursor to the beginning of the command line.
Ctrl W	Deletes the last word typed.
Esc B	Moves the cursor backward one word.
Esc D	Deletes from the cursor to the end of the word.
Esc F	Moves the cursor forward one word.
Delete or Backspace	Deletes the first character to the left of the cursor in the command line.

4

Using the HP Web Browser Interface

Contents

Overview

The HP web browser interface built into the switch lets you easily access the switch from a browser-based PC on your network. This lets you do the following:

- Optimize your network uptime by using the Alert Log and other diagnostic tools
- Make configuration changes to the switch
- Maintain security by configuring usernames and passwords

This chapter covers the following:

- General features (page 4-3).
- Starting a web browser interface session (page 4-4)
- Tasks for your first web browser interface session (page 4-7):
  - Creating usernames and passwords in the web browser interface (page 4-8)
  - Selecting the fault detection configuration for the Alert Log operation (page 4-23)
  - Getting access to online help for the web browser interface (page 4-11)
- Description of the web browser interface:
  - Overview window and tabs (page 4-15)
  - Port Utilization and Status displays (page 4-16)
  - Alert Log and Alert types (page 4-19)
  - Setting the Fault Detection Policy (page 4-23)

Note

If you want security beyond that achieved with user names and passwords, you can disable access to the web browser interface. This is done by either executing **no web-management** at the Command Prompt or changing the **Web Agent Enabled** parameter setting to **No (**page 6-4).

General Features

The Series 5300XL switch include these web browser interface features:

Switch Configuration:

- Ports
- VLANs and Primary VLAN
- Fault detection
- Port monitoring (mirroring)
- System information
- Enable/Disable Multicast Filtering (IGMP) and Spanning Tree
- IP
- Support and management URLs

Switch Security: Passwords

Switch Diagnostics:

- Ping/Link Test
- Device reset
- Configuration report

Switch status

- Port utilization
- Port counters
- Port status
- Alert log

Switch system information listing

Starting an HP Web Browser Interface Session with the Switch

You can start a web browser session in the following ways:

- Using a standalone web browser on a network connection from a PC or UNIX workstation:
  - Directly connected to your network
  - Connected through remote access to your network
- Using a management station running HP TopTools for Hubs & Switches on your network

Using a Standalone Web Browser in a PC or UNIX Workstation

This procedure assumes that you have a supported web browser installed on your PC or workstation, and that an IP address has been configured on the switch. (For more on assigning an IP address, refer to “IP Configuration” on page 7-3.)

1. Make sure the Java^{TM} applets are enabled for your browser. If they are not, use the options menu in your browser to do the following:
   - In Netscape, enable the **Java** and **JavaScript** options.
   - In Microsoft Internet Explorer, enable the **Java Permissions**.

Refer to your selected browser’s online Help for specific information on enabling the Java applets.

2. Type the IP address (or DNS name) of the switch in the browser **Location or Address** field and press [Enter]. (It is not necessary to include **http://**.)

   switch4108 [Enter] (example of a DNS-type name)

   10.11.12.195 [Enter] (example of an IP address)

   If you are using a Domain Name Server (DNS), your device may have a name associated with it (for example, **switch4108**) that you can type in the **Location or Address** field instead of the IP address. Using DNS names typically improves browser performance. See your network administrator for any name associated with the switch.

Using HP TopTools for Hubs & Switches

HP TopTools for Hubs & Switches is designed for installation on a network management workstation. For this reason, the HP TopTools system requirements are different from the system requirements for accessing the switch's web browser interface from a non-management PC or workstation. For HP TopTools requirements, refer to the information provided with HP TopTools for Hubs & Switches.

This procedure assumes that:

- You have installed the recommended web browser on a PC or workstation that serves as your network management station.
- The networked device you want to access has been assigned an IP address and (optionally) a DNS name and has been discovered by HP TopTools for Hubs & Switches. (For more on assigning an IP address, refer to "IP Configuration" on page 7-3.)

To establish a web browser session with HP TopTools running, do the following on the network management station:

1. Make sure the Java^{TM} applets are enabled for your web browser. If they are not, refer to the web browser online Help for specific information on enabling the Java applets.
2. Do *one* of the following tasks:
   - On the HP TopTools Maps view, double-click on the symbol for the networking device that you want to access.
   - In HP TopTools, in the Topology Information dialog box, in the device list, double-click on the entry for the device you want to access (IP address or DNS name).

3. The web browser interface automatically starts with the Status Overview window displayed for the selected device, as shown in figure 4-1.

Note

If the Registration window appears, click on the **Status** tab.

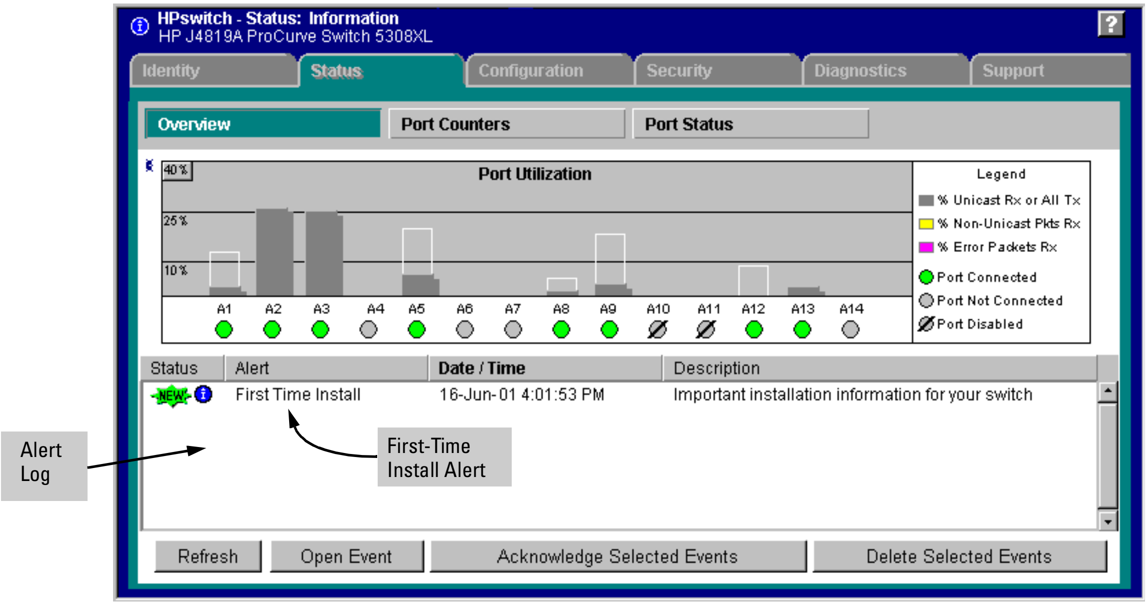

HPswitch - Status: Information
HP J4819A ProCurve Switch 5308XL

Status	Alert	Date / Time	Description
NEW	First Time Install	16-Jun-01 4:01:53 PM	Important installation information for your switch

Figure 4-1. Example of Status Overview Screen

Tasks for Your First HP Web Browser Interface Session

The first time you access the web browser interface, there are three tasks that you should perform:

- Review the "First Time Install" window
- Set Manager and Operator passwords
- Set access to the web browser interface online help

Viewing the "First Time Install" Window

When you access the switch's web browser interface for the first time, the Alert log contains a "First Time Install" alert, as shown in figure 4-2. This gives you information about first time installations, and provides an immediate opportunity to set passwords for security and to specify a Fault Detection policy, which determines the types of messages that will be displayed in the Alert Log.

Double click on **First Time Install** in the Alert log (figure 4-1 on page 4-6). The web browser interface then displays the "First Time Install" window, below.

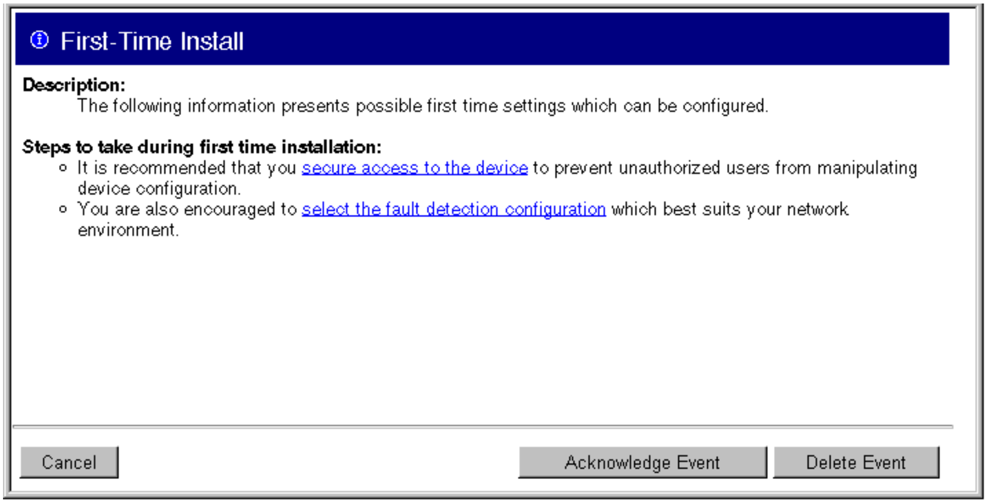
First-Time Install

Description:
The following information presents possible first time settings which can be configured.

Steps to take during first time installation:

- It is recommended that you secure access to the device to prevent unauthorized users from manipulating device configuration.
- You are also encouraged to select the fault detection configuration which best suits your network environment.

Figure 4-2. First-Time Install Window

This window is the launching point for the basic configuration you need to perform to set web browser interface passwords to maintain security and Fault Detection policy, which determines the types of messages that will be displayed in the Alert Log.

To set web browser interface passwords, click on **secure access to the device** to display the Device Passwords screen, and then go to the next page. (You can also access the password screen by clicking on the **Security** tab.)

To set Fault Detection policy, click on **select the fault detection configuration** in the second bullet in the window and go to the section, "Setting Fault Detection Policy" on page 4-23. (You can also access the password screen by clicking on the **Configuration** tab, and then [Fault Detection] button.)

Creating Usernames and Passwords in the Browser Interface

You may want to create both a username and password to create access security for your switch. There are two levels of access to the interface that can be controlled by setting user names and passwords:

- **Operator.** An Operator-level user name and password allows read-only access to most of the web browser interface, but prevents access to the Security window.
- **Manager.** A Manager-level user name and password allows full read/ write access to the web browser interface.

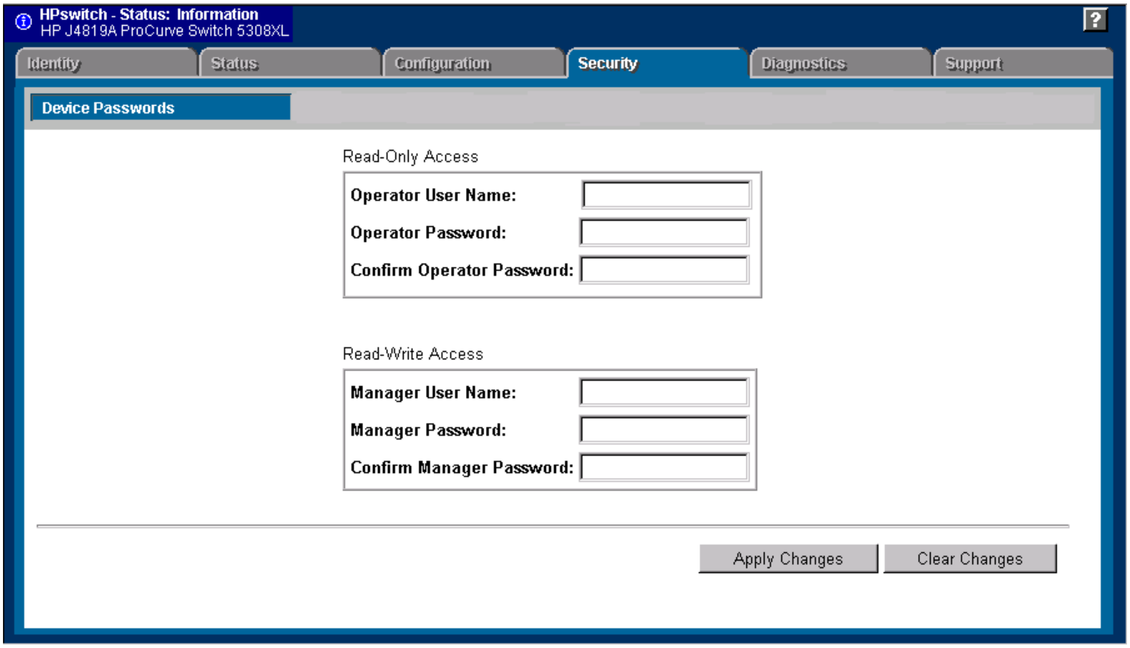

Figure 4-3. The Device Passwords Window

To set the passwords:

1. Access the Device Passwords screen by one of the following methods:

   - If the Alert Log includes a "First Time Install" event entry, double click on this event, then, in the resulting display, click on the **secure access to the device** link.

   - Select the **Security** tab.

2. Click in the appropriate box in the Device Passwords window and enter user names and passwords. You will be required to repeat the password strings in the confirmation boxes.

   Both the user names and passwords can be up to 16 printable ASCII characters.

3. Click on [Apply Changes] to activate the user names and passwords.

Note

Passwords you assign in the web browser interface will overwrite previous passwords assigned in either the web browser interface, the Command Prompt, or the switch console. That is, the most recently assigned passwords are the switch's passwords, regardless of which interface was used to assign the string.

Using the Passwords

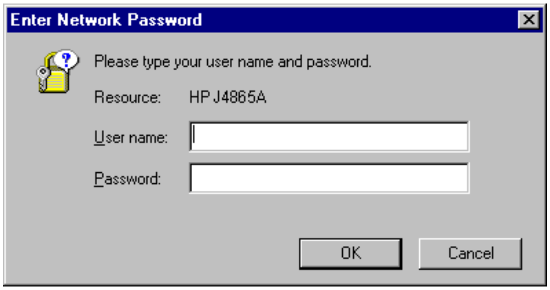

Please type your user name and password.

Resource: HP J4865A

Figure 4-4. Example of the Password Window in the Web Browser Interface

The manager and operator passwords are used to control access to all switch interfaces. Once set, you will be prompted to supply the password every time you try to access the switch through any of its interfaces. The password you enter determines the capability you have during that session:

- Entering the manager password gives you full read/write capabilities
- Entering the operator password gives you read and limited write capabilities.

Using the User Names

If you also set user names in the web browser interface screen, you must supply the correct user name for web browser interface access. If a user name has not been set, then leave the User Name field in the password window blank.

Note that the Command Prompt and switch console interfaces use only the password, and do not prompt you for the User Name.

If You Lose a Password

If you lose the passwords, you can clear them by pressing the Clear button on the front of the switch. *This action deletes all password and user name protection from all of the switch's interfaces.*

The Clear button is provided for your convenience, but its presence means that if you are concerned with the security of the switch configuration and operation, you should make sure the switch is installed in a secure location, such as a locked wiring closet.

Online Help for the HP Web Browser Interface

Online Help is available for the web browser interface. You can use it by clicking on the question mark button in the upper right corner of any of the web browser interface screens.

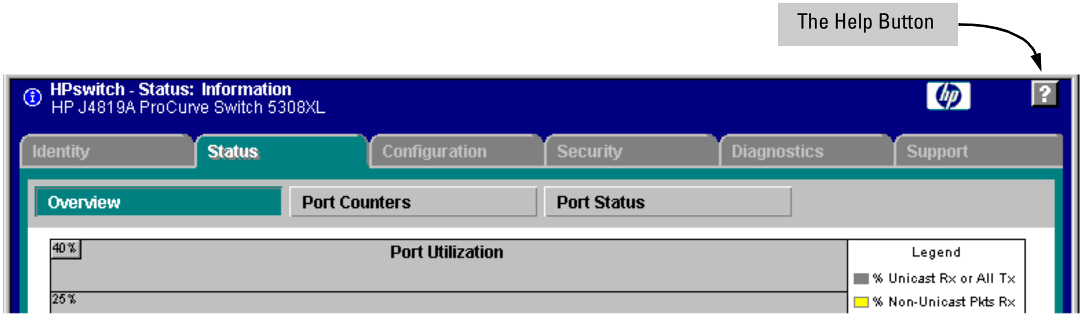
The Help Button

Figure 4-5. The Help Button

Context-sensitive help is provided for the screen you are on.

Note

If you do not have HP TopTools for Hubs and Switches installed on your network and do not have an active connection to the World Wide Web, then Online help for the web browser interface will not be available.

For more on Help access and operation, refer to "Help and the Management Server URL" on page 4-13.

Support/Mgmt URLs Feature

The Support/Mgmt URLs window enables you to change the World Wide Web Universal Resource Locator (URL) for two functions:

- **Support URL** – a support information site for your switch
- **Management Server URL** – the site for online help for the web browser interface, and, if set up, the URL of a network management station running HP TopTools for Hubs & Switches.

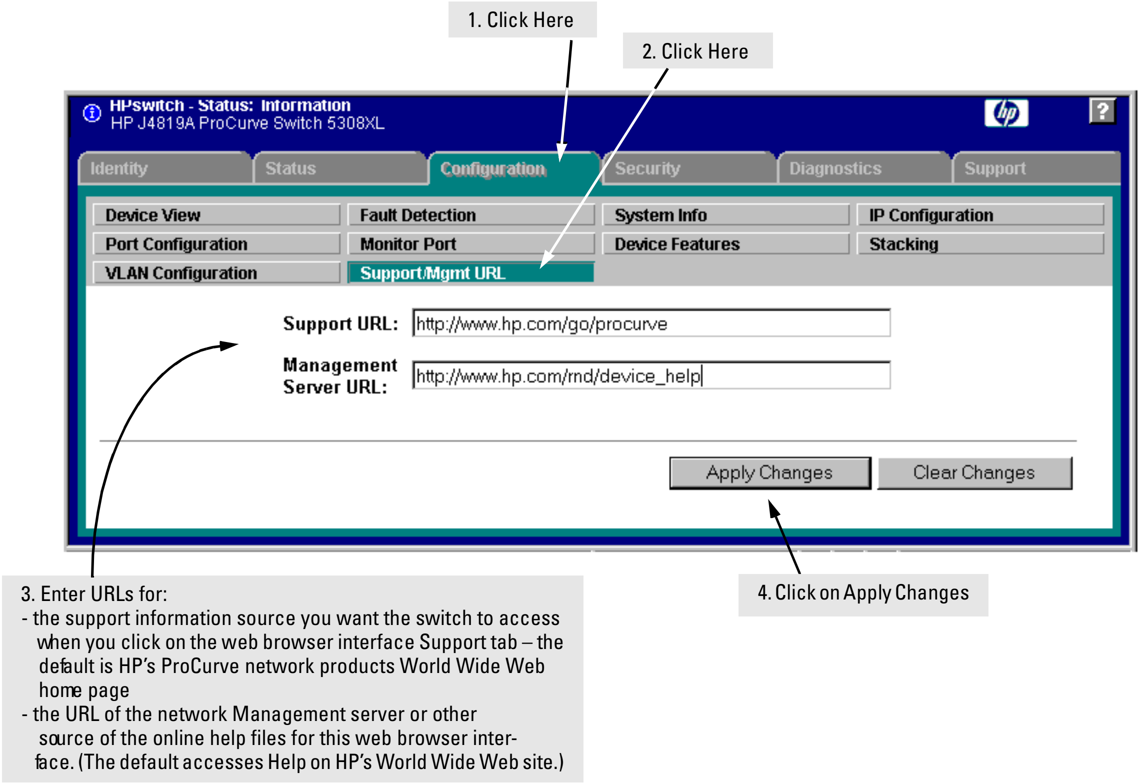

Support URL: http://www.hp.com/go/procurve

Management Server URL: http://www.hp.com/rnd/device_help

1. Click Here

2. Click Here

3. Enter URLs for:
- the support information source you want the switch to access when you click on the web browser interface Support tab – the default is HP's ProCurve network products World Wide Web home page
- the URL of the network Management server or other source of the online help files for this web browser interface. (The default accesses Help on HP's World Wide Web site.)

4. Click on Apply Changes

Figure 4-6. The Default Support/Mgmt URLs Window

Support URL

This is the site that the switch accesses when you click on the **Support** tab on the web browser interface. The default URL is:

http://www.hp.com/go/procurve

which is the World Wide Web site for Hewlett-Packard's networking products.

Click on the [Support] button on that page and you can get to support information regarding your switch, including white papers, operating system (OS) updates, and more.

You could instead enter the URL for a local site that you use for entering reports about network performance, or whatever other function you would like to be able to easily access by clicking on the [Support] tab.

Help and the Management Server URL

This field specifies which of the following two locations the switch will use to find online Help for the web browser interface:

- The URL of online Help provided by HP on the world wide web
- The URL of a network management station running HP TopTools for Hubs & Switches

Providing Online Help. *The Help files are automatically available if you install HP TopTools for Hubs & Switches on your network or if you already have Internet access to the World Wide Web.* (The Help files are included with HP TopTools for Hubs & Switches, and are also automatically available from HP via the World Wide Web.)

Retrieval of the Help files is controlled by automatic entries to the **Management Server URL** field on the **Configuration / Support/Mgmt URLs** screen, shown in figure 4-6. The switch is shipped with the URL set to retrieve online Help from the HP World Wide Web site. However, if HP TopTools for Hubs & Switches is installed on a management station on your network and discovers the switch, the Management Server URL is automatically changed to retrieve the Help from your TopTools management station.

If Online Help Fails To Operate. Do one of the following:

- If HP TopTools for Hubs & Switches is installed and running on your network, enter the IP address or DNS name of the network management station in the Management Server URL field shown in figure 4-7 on page 4-14.

- If you have World Wide Web access from your PC or workstation, and do not have HP TopTools installed on your network, enter the following URL in the Management Server URL field shown in figure 4-7 on page 4-14:

  http://www.hp.com/rnd/device_help

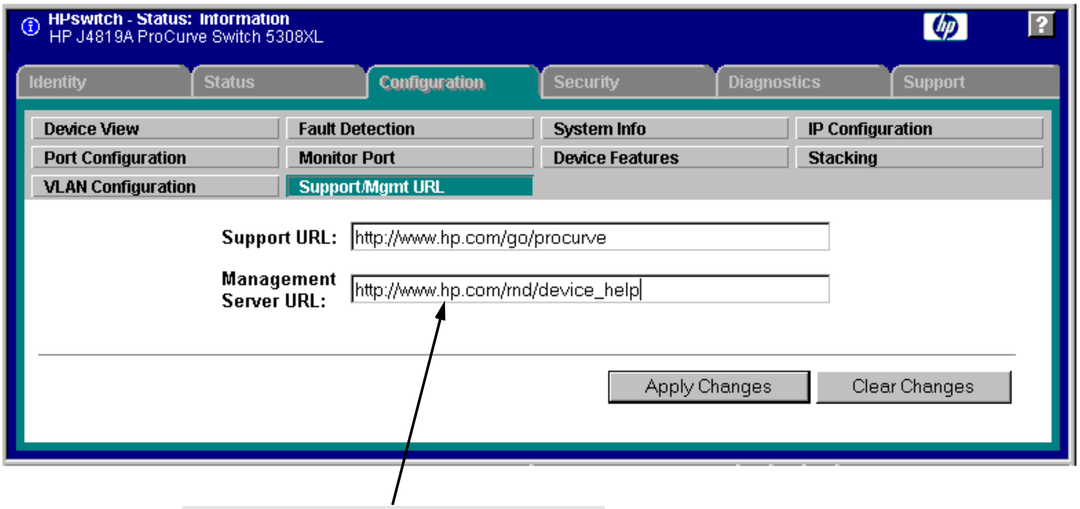

Support URL: http://www.hp.com/go/procurve

Management Server URL: http://www.hp.com/rnd/device_help

Enter IP address of HP TopTools network management station, or URL of location of help files on HP's World Wide Web site here.

Figure 4-7. How To Access Web Browser Interface Online Help

Policy Management and Configuration. HP Top Tools for Hubs & Switches can perform network-wide policy management and configuration of your switch. The Management Server URL field identifies the management station that is performing that function. For more information, refer to the documentation provided on the HP TopTools for Hubs & Switches CD shipped with the switch.

Status Reporting Features

Browser elements covered in this section include:

- The Overview window (below)
- Port utilization and status (page 4-16)
- The Alert log (page 4-19)
- The Status bar (page 4-22)

The Overview Window

The Overview Window is the home screen for any entry into the web browser interface.The following figure identifies the various parts of the screen.

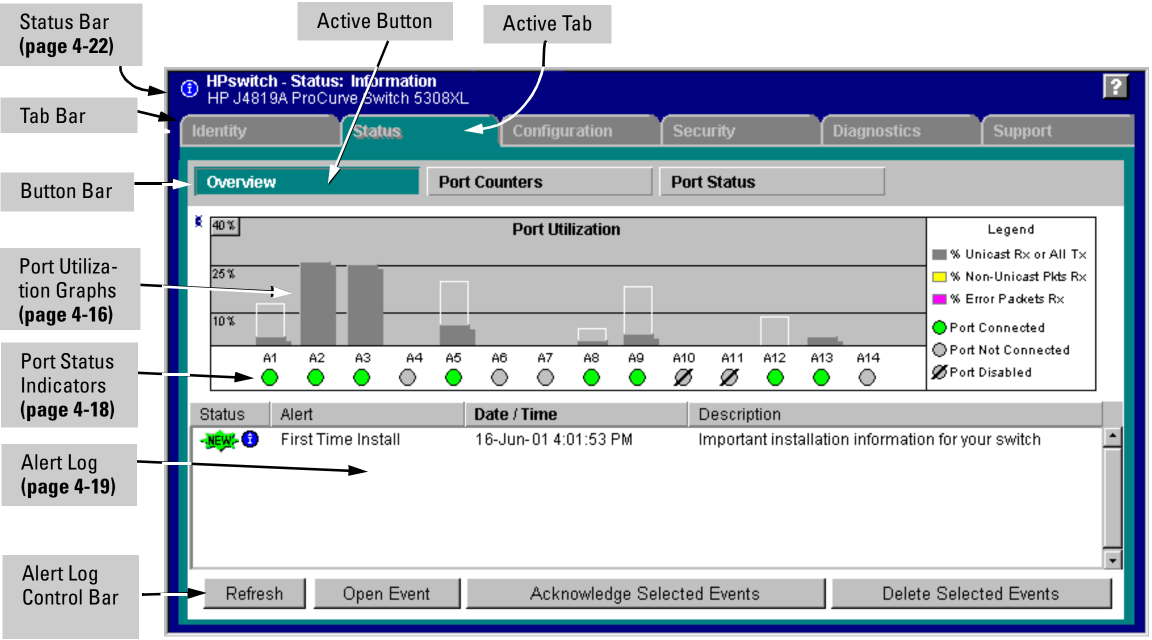

Status	Alert	Date / Time	Description
NEW	First Time Install	16-Jun-01 4:01:53 PM	Important installation information for your switch

Figure 4-8. The Overview Window

The Port Utilization and Status Displays

The Port Utilization and Status displays show an overview of the status of the switch and the amount of network activity on each port. The following figure shows a sample reading of the Port Utilization and Port Status.

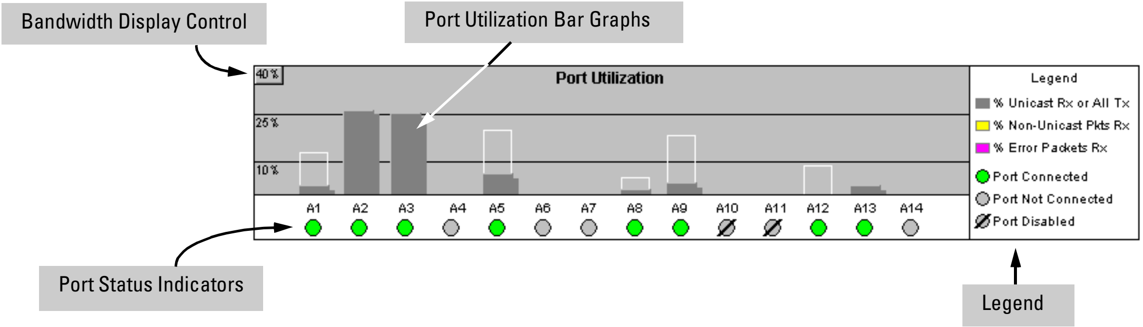

Bandwidth Display Control

Port Utilization Bar Graphs

Port Status Indicators

Legend

Figure 4-9. The Graphs Area

Port Utilization

The Port Utilization bar graphs show the network traffic on the port with a breakdown of the packet types that have been detected (unicast packets, non-unicast packets, and error packets). The Legend identifies traffic types and their associated colors on the bar graph:

- **% Unicast Rx & All Tx:** This is all unicast traffic received and all transmitted traffic of any type. This indicator (a blue color on many systems) can signify either transmitted or received traffic.
- **% Non-Unicast Pkts Rx:** All multicast and broadcast traffic received by the port. This indicator (a gold color on many systems) enables you to know “at-a-glance” the source of any non-unicast traffic that is causing high utilization of the switch. For example, if one port is receiving heavy broadcast or multicast traffic, all ports will become highly utilized. By color-coding the received broadcast and multicast utilization, the bar graph quickly and easily identifies the offending port. This makes it faster and easier to discover the exact source of the heavy traffic because you don’t have to examine port counter data from several ports.
- **% Error Pkts Rx**: All error packets received by the port. (This indicator is a reddish color on many systems.) Although errors received on a port are not propagated to the rest of the network, a consistently high number of errors on a specific port may indicate a problem on the device or network segment connected to the indicated port.

- **Maximum Activity Indicator:** As the bars in the graph area change height to reflect the level of network activity on the corresponding port, they leave an outline to identify the maximum activity level that has been observed on the port.

Utilization Guideline. A network utilization of 40% is considered the maximum that a typical Ethernet-type network can experience before encountering performance difficulties. If you observe utilization that is consistently higher than 40% on any port, click on the Port Counters button to get a detailed set of counters for the port.

To change the amount of bandwidth the Port Utilization bar graph shows. Click on the bandwidth display control button in the upper left corner of the graph. (The button shows the current scale setting, such as 40%.) In the resulting menu, select the bandwidth scale you want the graph to show (3%, 10%, 25%, 40%, 75%, or 100%), as shown in figure figure 4-10.

Note that when viewing activity on a gigabit port, you may want to select a lower value (such as 3% or 10%). This is because the bandwidth utilization of current network applications on gigabit links is typically minimal, and may not appear on the graph if the scale is set to show high bandwidth utilization.

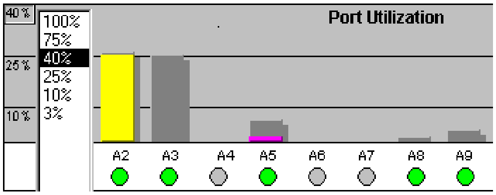

Figure 4-10. Changing the Graph Area Scale

To display values for each graph bar. Hold the mouse cursor over any of the bars in the graph, and a pop-up display is activated showing the port identification and numerical values for each of the sections of the bar, as shown in figure 4-11 (next).

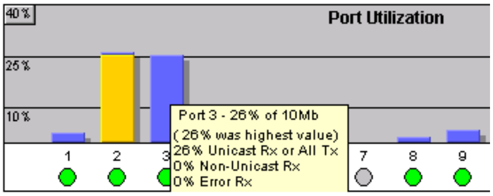

Port 3 - 26% of 10Mb
(26% was highest value)
26% Unicast Rx or All Tx
0% Non-Unicast Rx
0% Error Rx

Figure 4-11. Display of Numerical Values for the Bar

Port Status

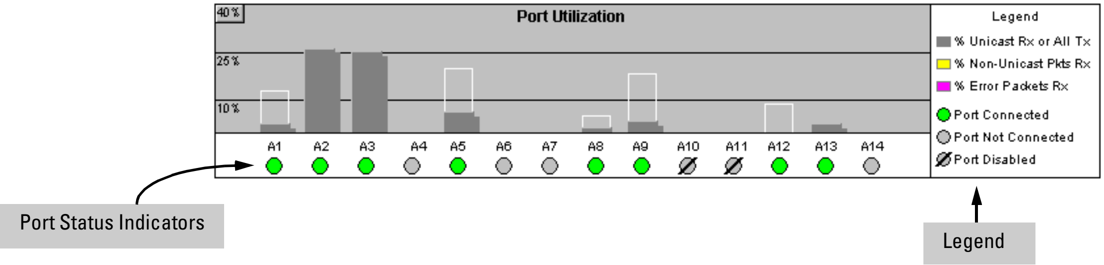

Port Status Indicators

Legend

Figure 4-12. The Port Status Indicators and Legend

The Port Status indicators show a symbol for each port that indicates the general status of the port. There are four possible statuses:

- **Port Connected** – the port is enabled and is properly connected to an active network device.
- **Port Not Connected** – the port is enabled but is not connected to an active network device. A cable may not be connected to the port, or the device at the other end may be powered off or inoperable, or the cable or connected device could be faulty.
- **Port Disabled** – the port has been configured as disabled through the web browser interface, the switch console, or SNMP network management.
- **Port Fault-Disabled** – a fault condition has occurred on the port that has caused it to be auto-disabled. Note that the Port Fault-Disabled symbol will be displayed in the legend only if one or more of the ports is in that status. See appendix B, “Monitoring and Analyzing Switch Operation” for more information.

The Alert Log

The web browser interface Alert Log, shown in the lower half of the screen, shows a list of network occurrences, or *alerts*, that were detected by the switch. Typical alerts are **Broadcast Storm**, indicating an excessive number of broadcasts received on a port, and **Problem Cable**, indicating a faulty cable. A full list of alerts is shown in the table on page 4-20.

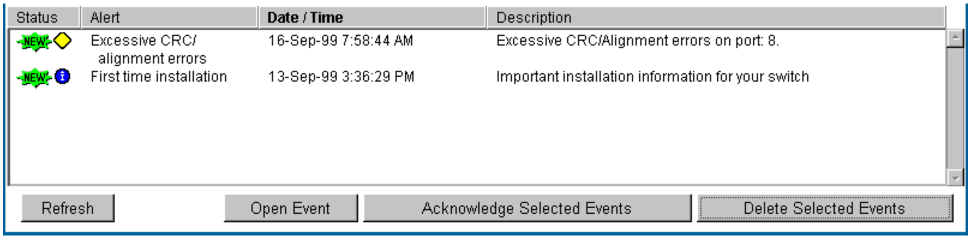

Status	Alert	**Date / Time**	Description
NEW ◇	Excessive CRC/ alignment errors	16-Sep-99 7:58:44 AM	Excessive CRC/Alignment errors on port: 8.
NEW ℹ	First time installation	13-Sep-99 3:36:29 PM	Important installation information for your switch

Figure 4-13. Example of the Alert Log

Each alert has the following fields of information:

- **Status** – The level of severity of the event generated. Severity levels can be Information, Normal, Warning, and Critical. If the alert is new (has not yet been acknowledged), the New symbol is also in the Status column.
- **Alert** – The specific event identification.
- **Date/Time** – The date and time the event was received by the web browser interface. This value is shown in the format: ***DD-MM-YY HH:MM:SS*** **AM/PM**, for example, **16-Sep-99 7:58:44 AM**.
- **Description** – A short narrative statement that describes the event. For example, **Excessive CRC/Alignment errors on port: 8**.

Sorting the Alert Log Entries

The alerts are sorted, by default, by the Date/Time field with the most recent alert listed at the top of the list. The second most recent alert is displayed below the top alert and so on. If alerts occurred at the same time, the simultaneous alerts are sorted by order in which they appear in the MIB.

The alert field that is being used to sort the alert log is indicated by which column heading is in bold. You can sort by any of the other columns by clicking on the column heading. The Alert and Description columns are sorted alphabetically, while the Status column is sorted by severity type, with more critical severity indicators appearing above less critical indicators.

Alert Types

The following table lists the types of alerts that can be generated.

Table 4-1. Alert Strings and Descriptions

Alert String	**Alert Description**
First Time Install	Important installation information for your switch.
Too many undersized/ giant packets	A device connected to this port is transmitting packets shorter than 64 bytes or longer than 1518 bytes (longer than 1522 bytes if tagged), with valid CRCs (unlike runts, which have invalid CRCs).
Excessive jabbering	A device connected to this port is incessantly transmitting packets ("jabbering"), detected as oversized packets with CRC errors.
Excessive CRC/alignment errors	A high percentage of data errors has been detected on this port. Possible causes include: • Faulty cabling or invalid topology. • Duplex mismatch (full-duplex configured on one end of the link, half-duplex configured on the other) • A malfunctioning NIC, NIC driver, or transceiver
Excessive late collisions	Late collisions (collisions detected after transmitting 64 bytes) have been detected on this port. Possible causes include: • An overextended LAN topology • Duplex mismatch (full-duplex configured on one end of the link, half-duplex configured on the other) • A misconfigured or faulty device connected to the port
High collision or drop rate	A large number of collisions or packet drops have occurred on the port. Possible causes include: • A extremely high level of traffic on the port • Duplex mismatch • A misconfigured or malfunctioning NIC or transceiver on a device connected to this port • A topology loop in the network
Excessive broadcasts	An extremely high percentage of broadcasts was received on this port. This degrades the performance of all devices connected to the port. Possible causes include: • A network topology loop—this is the usual cause • A malfunctioning device, NIC, NIC driver, or software package
Loss of Link	Lost connection to one or multiple devices on the port.

Note

When troubleshooting the sources of alerts, it may be helpful to check the switch's Port Status and Port Counter windows and the Event Log in the console interface.

Viewing Detail Views of Alert Log Entries

By double clicking on Alert Entries, the web browser interface displays a Detail View or separate window detailing information about the events. The Detail View contains a description of the problem and a possible solution. It also provides four management buttons:

- **Acknowledge Event** – removes the New symbol from the log entry
- **Delete Event** – removes the alert from the Alert Log
- **Cancel Button** – closes the detail view with no change to the status of the alert and returns you to the Overview screen.

A sample Detail View describing an Excessive CRC/Alignment Error alert is shown here.

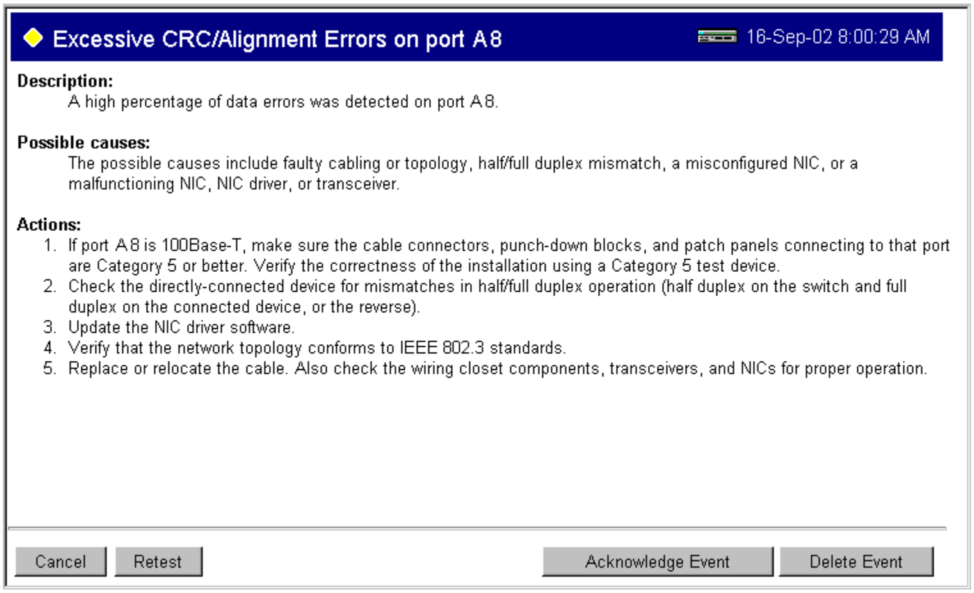

◆ Excessive CRC/Alignment Errors on port A8 16-Sep-02 8:00:29 AM

Description:
A high percentage of data errors was detected on port A8.

Possible causes:
The possible causes include faulty cabling or topology, half/full duplex mismatch, a misconfigured NIC, or a malfunctioning NIC, NIC driver, or transceiver.

Actions:
1. If port A8 is 100Base-T, make sure the cable connectors, punch-down blocks, and patch panels connecting to that port are Category 5 or better. Verify the correctness of the installation using a Category 5 test device.
2. Check the directly-connected device for mismatches in half/full duplex operation (half duplex on the switch and full duplex on the connected device, or the reverse).
3. Update the NIC driver software.
4. Verify that the network topology conforms to IEEE 802.3 standards.
5. Replace or relocate the cable. Also check the wiring closet components, transceivers, and NICs for proper operation.

Figure 4-14. Example of Alert Log Detail View

The Status Bar

The Status Bar is displayed in the upper left corner of the web browser interface screen. Figure 4-15 shows an expanded view of the status bar.

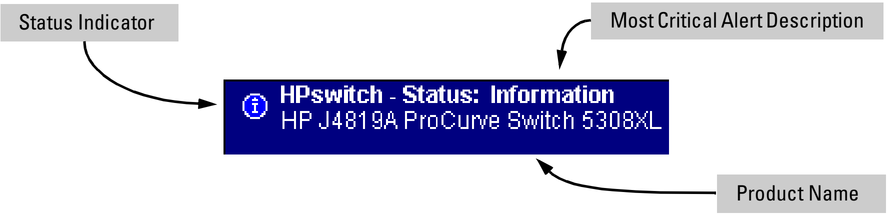

HPswitch - Status: Information
HP J4819A ProCurve Switch 5308XL

Figure 4-15. Example of the Status Bar

The Status bar consists of four objects:

- **Status Indicator.** Indicates, by icon, the severity of the most critical alert in the current display of the Alert Log. This indicator can be one of three shapes and colors as shown in the following table.

Table 4-2. Status Indicator Key

Color	Switch Status	Status Indicator Shape
Blue	Normal Activity; "First time installation" information available in the Alert log.	ⓘ
Green	Normal Activity	●
Yellow	Warning	◆
Red	Critical	⚠

- **System Name.** The name you have configured for the switch by using Identity screen, **system name** command, or the switch console **System Information** screen.
- **Most Critical Alert Description.** A brief description of the earliest, unacknowledged alert with the current highest severity in the Alert Log, appearing in the right portion of the Status Bar. In instances where multiple critical alerts have the same severity level, only the earliest unacknowledged alert is deployed in the Status bar.

- **Product Name.** The product name of the switch to which you are connected in the current web browser interface session.

Setting Fault Detection Policy

One of the powerful features in the web browser interface is the Fault Detection facility. For your switch, this feature controls the types of alerts reported to the Alert Log based on their level of severity.

Set this policy in the Fault Detection window (figure 4-16).

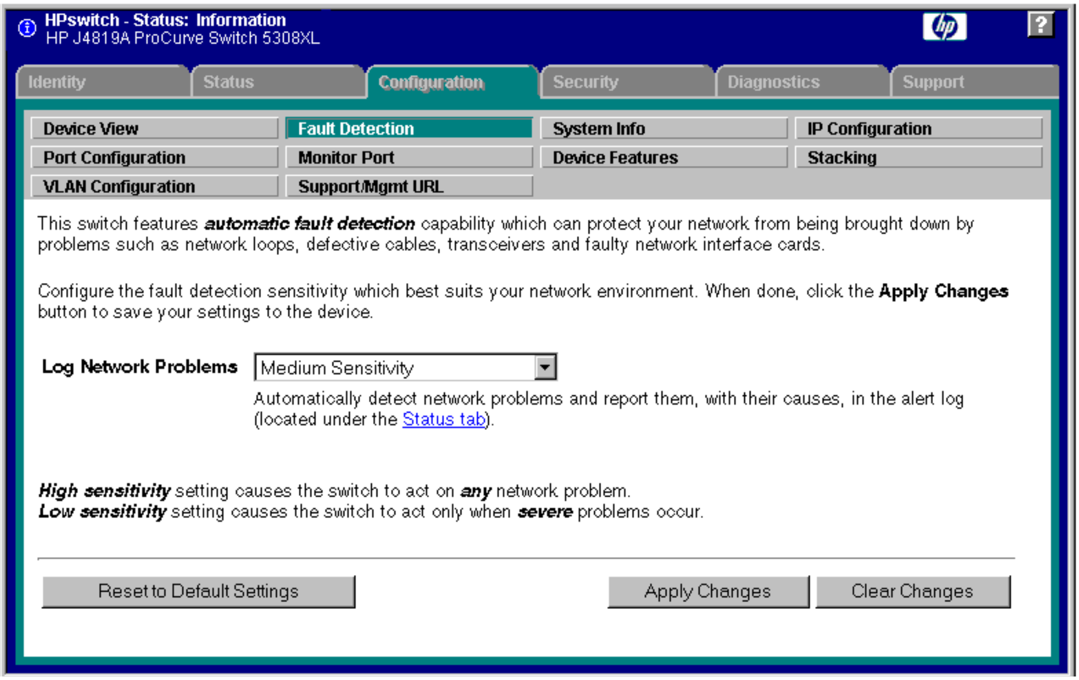

This switch features ***automatic fault detection*** capability which can protect your network from being brought down by problems such as network loops, defective cables, transceivers and faulty network interface cards.

Configure the fault detection sensitivity which best suits your network environment. When done, click the **Apply Changes** button to save your settings to the device.

Log Network Problems Medium Sensitivity

Automatically detect network problems and report them, with their causes, in the alert log (located under the Status tab).

High sensitivity setting causes the switch to act on ***any*** network problem.
Low sensitivity setting causes the switch to act only when ***severe*** problems occur.

Figure 4-16. The Fault Detection Window

The Fault Detection screen contains a list box for setting fault detection and response policy. You set the sensitivity level at which a network problem should generate an alert and send it to the Alert Log.

To provide the most information on network problems in the Alert Log, the recommended sensitivity level for **Log Network Problems** is **High Sensitivity**. The Fault Detection settings are:

- **High Sensitivity.** This policy directs the switch to send all alerts to the Alert Log. This setting is most effective on networks that have none or few problems.
- **Medium Sensitivity.** This policy directs the switch to send alerts related to network problems to the Alert Log. If you want to be notified of problems which cause a noticeable slowdown on the network, use this setting.
- **Low Sensitivity.** This policy directs the switch to send only the most severe alerts to the Alert Log. This policy is most effective on a network that normally has a lot of problems and you want to be informed of only the most severe ones.
- **Never.** Disables the Alert Log and transmission of alerts (traps) to the management server (in cases where a network management tool such as HP TopTools for Hubs & Switches is in use). Use this option when you don't want to use the Alert Log.

The Fault Detection Window also contains three Change Control Buttons:

- **Apply Changes.** This button stores the settings you have selected for all future sessions with the web browser interface until you decide to change them.
- **Clear Changes.** This button removes your settings and returns the settings for the list box to the level it was at in the last saved detection-setting session.
- **Reset to Default Settings.** This button reverts the policy setting to Medium Sensitivity for Log Network Problems.

Switch Memory and Configuration

Contents

Overview

This chapter describes:

- How switch memory manages configuration changes
- How the CLI implements configuration changes
- How the menu interface and web browser interface implement configuration changes
- How the switch provides OS (operating system) options through primary/ secondary flash image options
- How to use the switch's primary and secondary flash options, including displaying flash information, booting or restarting the switch, and other topics

Overview of Configuration File Management

The switch maintains two configuration files, the *running-config* file and the *startup-config* file.

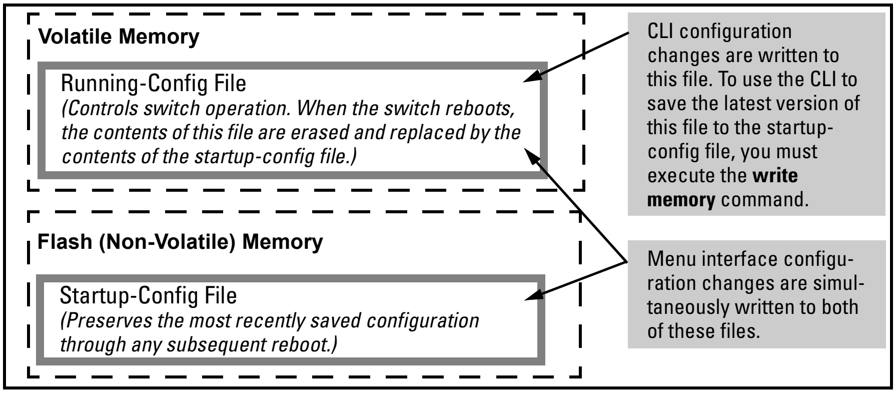
Volatile Memory
Running-Config File
(Controls switch operation. When the switch reboots, the contents of this file are erased and replaced by the contents of the startup-config file.)
Flash (Non-Volatile) Memory
Startup-Config File
(Preserves the most recently saved configuration through any subsequent reboot.)
CLI configuration changes are written to this file. To use the CLI to save the latest version of this file to the startup-config file, you must execute the **write memory** command.
Menu interface configuration changes are simultaneously written to both of these files.

Figure 5-1. Conceptual Illustration of Switch Memory Operation

- **Running Config File:** Exists in volatile memory and controls switch operation. If no configuration changes have been made in the CLI since the switch was last booted, the running-config file is identical to the startup-config file.
- **Startup-config File:** Exists in flash (non-volatile) memory and is used to preserve the most recently-saved configuration as the "permanent" configuration.

Rebooting the switch replaces the current running-config file with a new running-config file that is an exact copy of the current startup-config file.

Note

Any of the following actions reboots the switch:

- Executing the **boot** or the **reload** command in the CLI
- Executing the **Reboot** command in the menu interface
- Pressing the Reset button on the front of the switch
- Removing, then restoring power to the switch

For more on reboots and the switch's dual-flash images, see "Using Primary and Secondary Flash Image Options" on page 5-12.

Options for Saving a New Configuration. Making one or more changes to the running-config file creates a new operating configuration. *Saving* a new configuration means to overwrite (replace) the current startup-config file with the current running-config file. This means that if the switch subsequently reboots for any reason, it will resume operation using the new configuration instead of the configuration previously defined in the startup-config file. There are three ways to save a new configuration:

- **In the CLI:** Use the **write memory** command. This overwrites the current startup-config file with the contents of the current running-config file.
- **In the menu interface:** Use the **Save** command. This overwrites *both* the running-config file and the startup-config file with the changes you have specified in the menu interface screen.
- **In the web browser interface:** Use the [Apply Changes] button or other appropriate button. This overwrites *both* the running-config file and the startup-config file with the changes you have specified in the web browser interface window.

Note that using the CLI instead of the menu or web browser interface gives you the option of changing the running configuration without affecting the startup configuration. This allows you to test the change without making it "permanent". When you are satisfied that the change is satisfactory, you can make it permanent by executing the **write memory** command. For example, suppose you use the following command to disable port 5:

```
HPswitch(config)# interface ethernet 5 disable
```

The above command disables port 5 in the running-config file, but not in the startup-config file. Port 5 remains disabled only until the switch reboots. If you want port 5 to remain disabled through the next reboot, use **write memory** to save the current running-config file to the startup-config file in flash memory.

```
HPswitch(config)# write memory
```

If you use the CLI to make a configuration change and then change from the CLI to the Menu interface without first using write memory to save the change to the startup-config file, then the switch prompts you to save the change. For example, if you use the CLI to create VLAN 20, and then select the menu interface, VLAN 20 is configured in the running-config file, but not in the startup-config file. In this case you will see:

```
HPswitch(config)# vlan 20
HPswitch(config)# menu
Do you want to save current configuration [y/n]?
```

If you type [Y], the switch overwrites the startup-config file with the running-config file, and your configuration change(s) will be preserved across reboots. If you type [N], your configuration change(s) will remain only in the running-config file. In this case, if you do not subsequently save the running-config file, your unsaved configuration changes will be lost if the switch reboots for any reason.

Storing and Retrieving Configuration Files. You can store or retrieve a backup copy of the startup-config file on another device. For more information, see appendix A, "Transferring an Operating System or Startup-Config File"

Using the CLI To Implement Configuration Changes

The CLI offers these capabilities:

- Access to the full set of switch configuration features
- The option of testing configuration changes before making them permanent

How To Use the CLI To View the Current Configuration Files. Use **show** commands to view the configuration for individual features, such as port status or Spanning Tree Protocol. However, to view either the entire startup-config file or the entire running-config file, use the following commands:

- **show config** — Displays a listing of the current startup-config file.
- **show running-config —** Displays a listing of the current running-config file.
- **write terminal** — Displays a listing of the current running-config file.
- **show config status** — Compares the startup-config file to the running-config file and lists one of the following results:
  - If the two configurations are the same you will see:
    - `Running configuration is the same as the startup configuration.`
  - If the two configurations are different, you will see:
    - `Running configuration has been changed and needs to be saved.`

Note

Show config, **show running-config**, and **write terminal** commands display the configuration settings that differ from the switch's factory-default configuration.

How To Use the CLI To Reconfigure Switch Features. Use this procedure to permanently change the switch configuration (that is, to enter a change in the startup-config file).

1. Use the appropriate CLI commands to reconfigure the desired switch parameters. This updates the selected parameters in the running-config file.
2. Use the appropriate **show** commands to verify that you have correctly made the desired changes.

3. Observe the switch's performance with the new parameter settings to verify the effect of your changes.
4. When you are satisfied that you have the correct parameter settings, use the **write memory** command to copy the changes to the startup-config file.

Syntax: `write memory`

For example, the default port mode setting is **auto**. Suppose that your network uses Cat 3 wiring and you want to connect the switch to another autosensing device capable of 100 Mbps operation. Because 100 Mbps over Cat 3 wiring can introduce transmission problems, the recommended port mode is **auto-10**, which allows the port to negotiate full- or half-duplex, but restricts speed to 10 Mbps. The following command configures port A5 to auto-10 mode in the running-config file, allowing you to observe performance on the link without making the mode change permanent.

```
HPswitch(config)# interface e a5 speed-duplex auto-10
```

After you are satisfied that the link is operating properly, you can save the change to the switch's permanent configuration (the startup-config file) by executing the following command:

```
HPswitch(config)# write memory
```

The new mode (**auto-10**) on port A5 is now saved in the startup-config file, and the startup-config and running-config files are identical. If you subsequently reboot the switch, the **auto-10** mode configuration on port A5 will remain because it is included in the startup-config file.

How To Cancel Changes You Have Made to the Running-Config File.

If you use the CLI to change parameter settings in the running-config file, and then decide that you don't want those changes to remain, you can use either of the following methods to remove them:

- Manually enter the earlier values you had for the changed settings. (This is recommended if you want to restore a small number of parameter settings to their previous boot-up values.)
- Update the running-config file to match the startup-config file by rebooting the switch. (This is recommended if you want to restore a larger number of parameter settings to their previous boot-up values.)

If you use the CLI to change a parameter setting, and then execute the **boot** command without first executing the **write memory** command to save the change, the switch prompts you to specify whether to save the changes in the current running-config file. For example:

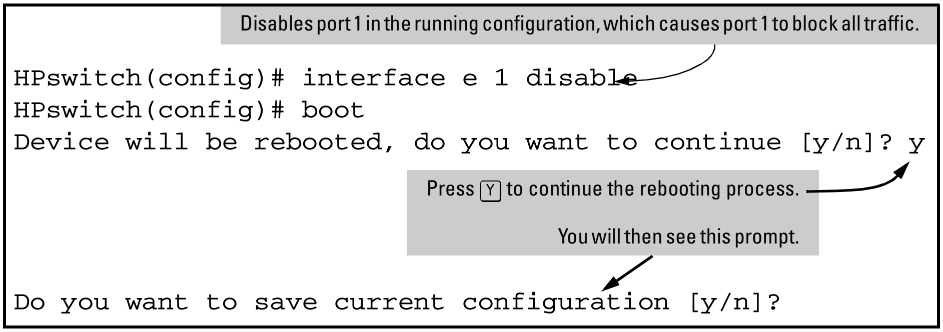

Disables port 1 in the running configuration, which causes port 1 to block all traffic.

```
HPswitch(config)# interface e 1 disable
HPswitch(config)# boot
Device will be rebooted, do you want to continue [y/n]? y
```

Press [Y] to continue the rebooting process.

You will then see this prompt.

```
Do you want to save current configuration [y/n]?
```

Figure 5-2. Boot Prompt for an Unsaved Configuration

The above prompt means that one or more parameter settings in the running-config file differ from their counterparts in the startup-config file and you need to choose which config file to retain and which to discard.

- If you want to update the startup-config file to match the running-config file, press [Y] for “yes”. (This means that the changes you entered in the running-config file will be saved in the startup-config file.)
- If you want to discard the changes you made to the running-config file so that it will match the startup-config file, then press [N] for “no”. (This means that the switch will discard the changes you entered in the running-config file and will update the running-config file to match the startup-config file.)

Note

If you use the CLI to make a change to the running-config file, you should either use the **write memory** command or select the save option allowed during a reboot (figure 5-5-2, above) to save the change to the startup-config file. That is, if you use the CLI to change a parameter setting, but then reboot the switch from either the CLI or the menu interface without first executing the **write memory** command in the CLI, the current startup-config file will replace the running-config file, and any changes in the running-config file will be lost.

Using the **Save** command in the menu interface does not save a change made to the running config by the CLI unless you have also made a configuration change in the menu interface. Also, the menu interface displays the current running-config values. Thus, where a parameter setting is accessible from both the CLI and the menu interface, if you change the setting in the CLI, the new value will appear in the menu interface display for that parameter. *However, as indicated above, unless you also make a configuration change in the menu interface, only the* **write memory** *command in the CLI will actually save the change to the startup-config file.*

How To Reset the startup-config and running-config Files to the Factory Default Configuration. This command reboots the switch, replacing the contents of the current startup-config and running-config files with the factory-default startup configuration.

Syntax: `erase startup-config`

For example:

```
HPswitch(config)# erase startup-config
Configuration will be deleted and device rebooted, continue [y/n]?
```

Press Y to replace the current configuration with the factory default configuration and reboot the switch. Press N to retain the current configuration and prevent a reboot.

Using the Menu and Web Browser Interfaces To Implement Configuration Changes

The menu and web browser interfaces offer these advantages:

- Quick, easy menu or window access to a subset of switch configuration features
- Viewing several related configuration parameters in the same screen, with their default and current settings
- Immediately changing both the running-config file and the startup-config file with a single command

Menu: Implementing Configuration Changes

You can use the menu interface to simultaneously save and implement a subset of switch configuration changes without having to reboot the switch. That is, when you save a configuration change in the menu interface, you simultaneously change both the running-config file and the startup-config file.

Note

The only exception to this operation are two VLAN-related parameter changes that require a reboot—described under "Rebooting To Activate Configuration Changes" on page 5-10.

Using **Save** and **Cancel** in the Menu Interface

For any configuration screen in the menu interface, the Save command:

1. Implements the changes in the running-config file
2. Saves your changes to the startup-config file

If you decide not to save and implement the changes in the screen, select **Cancel** to discard them and continue switch operation with the current operation. For example, suppose you have made the changes shown below in the System Information screen:

To save and implement the changes for all parameters in this screen, press the [Enter] key, then press [S] (for **Save**). To cancel all changes, press the [Enter] key, then press [C] (for **Cancel**)

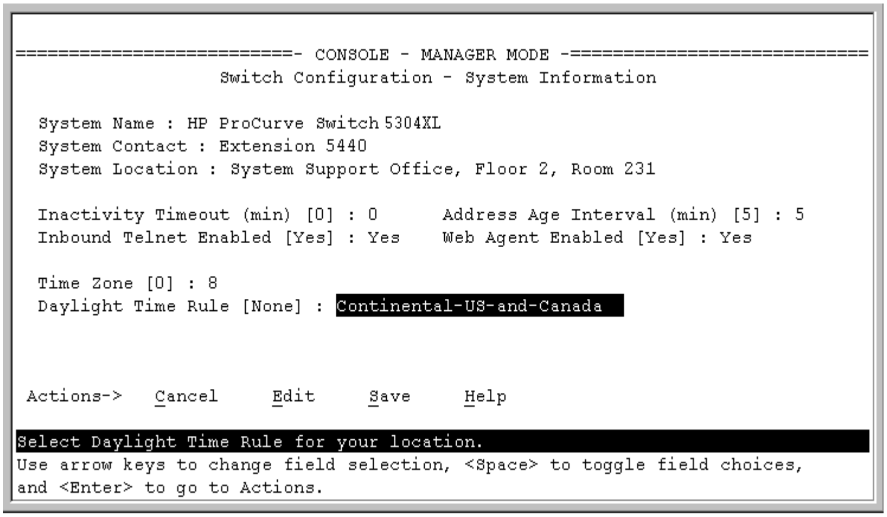

```
==========================- CONSOLE - MANAGER MODE -============================
                   Switch Configuration - System Information

  System Name : HP ProCurve Switch 5304XL
  System Contact : Extension 5440
  System Location : System Support Office, Floor 2, Room 231

  Inactivity Timeout (min) [0] : 0      Address Age Interval (min) [5] : 5
  Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes

  Time Zone [0] : 8
  Daylight Time Rule [None] : Continental-US-and-Canada


 Actions->   Cancel     Edit     Save     Help

Select Daylight Time Rule for your location.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Figure 5-3. Example of Pending Configuration Changes that Can Be Saved or Cancelled

Note

If you reconfigure a parameter in the CLI and then go to the menu interface without executing a **write memory** command, those changes are stored only in the running configuration (even if you execute a Save operation in the menu interface). If you then execute a switch reboot command in the menu interface, the switch discards the configuration changes made while using the CLI. To ensure that changes made while using the CLI are saved, execute **write memory** in the CLI before rebooting the switch.

Rebooting from the Menu Interface

- Terminates the current session and performs a reset of the operating system
- Activates any configuration changes that require a reboot
- Resets statistical counters to zero

(Note that statistical counters can be reset to zero without rebooting the switch. See “Displaying Port Counters” on page B-11.)

To Reboot the switch, use the **Reboot Switch** option in the Main Menu. (Note that the Reboot Switch option is not available if you log on in Operator mode; that is, if you enter an Operator password instead of a manager password at the password prompt.)

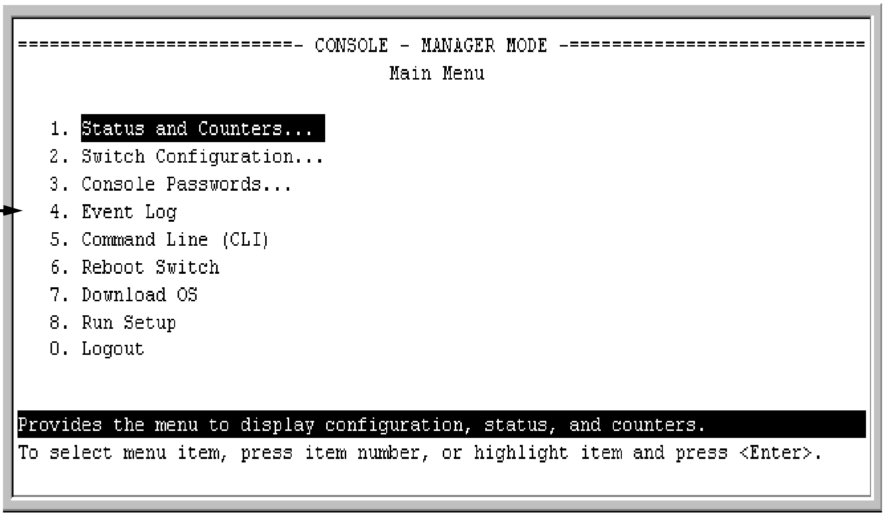

```
=========================- CONSOLE - MANAGER MODE -============================
                                  Main Menu

   1. Status and Counters...
   2. Switch Configuration...
   3. Console Passwords...
   4. Event Log
   5. Command Line (CLI)
   6. Reboot Switch
   7. Download OS
   8. Run Setup
   0. Logout


Provides the menu to display configuration, status, and counters.
To select menu item, press item number, or highlight item and press <Enter>.
```

Reboot Switch option

Figure 5-4. The Reboot Switch Option in the Main Menu

Rebooting To Activate Configuration Changes. Configuration changes for most parameters become effective as soon as you save them. However, you must reboot the switch in order to implement a change in the **Maximum VLANs to support** parameter.

(To access these parameters, go to the Main menu and select **2. Switch Configuration**, then **8. VLAN Menu**, then **1. VLAN Support**.)

If configuration changes requiring a reboot have been made, the switch displays an asterisk (*) next to the menu item in which the change has been made. For example, if you change and save parameter values for the **Maximum VLANs to support** parameter, an asterisk appears next to the **VLAN Support** entry in the VLAN Menu screen, and also next to the **Switch Configuration** . . . entry in the Main menu, as shown in figure 4-6:

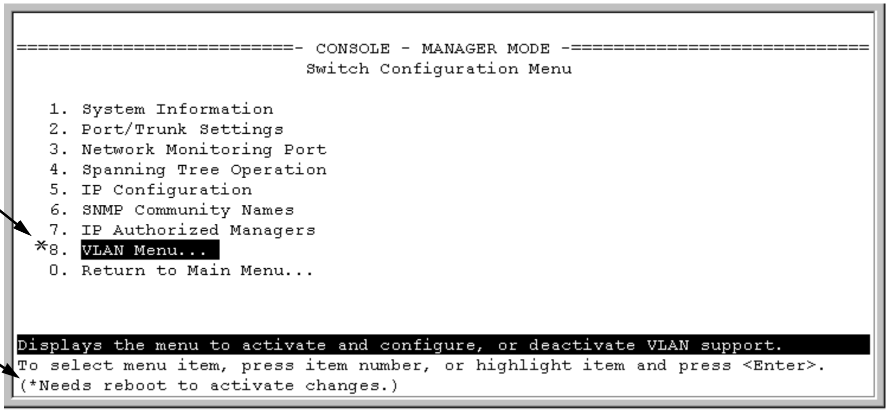

```
=========================- CONSOLE - MANAGER MODE -===========================
                          Switch Configuration Menu

   1. System Information
   2. Port/Trunk Settings
   3. Network Monitoring Port
   4. Spanning Tree Operation
   5. IP Configuration
   6. SNMP Community Names
   7. IP Authorized Managers
  *8. VLAN Menu...
   0. Return to Main Menu...


Displays the menu to activate and configure, or deactivate VLAN support.
To select menu item, press item number, or highlight item and press <Enter>.
(*Needs reboot to activate changes.)
```

Asterisk indicates a configuration change that requires a reboot in order to take effect.

Reminder to reboot the switch to activate configuration changes.

Figure 5-5. Indication of a Configuration Change Requiring a Reboot

Web: Implementing Configuration Changes

You can use the web browser interface to simultaneously save and implement a subset of switch configuration changes without having to reboot the switch. That is, when you save a configuration change (in most cases, by clicking on Apply Changes or Apply Settings, you simultaneously change both the running-config file and the startup-config file.

Note

If you reconfigure a parameter in the CLI and then go to the browser interface without executing a **write memory** command, those changes will be saved to the startup-config file if you click on Apply Changes or Apply Settings in the web browser interface.

Using Primary and Secondary Flash Image Options

The Series 4100GL switches feature two flash memory locations for storing system image (operating system, or OS) files:

- **Primary Flash:** The default storage for OS (system image) files.
- **Secondary Flash:** The additional storage for either a redundant or an alternate OS (system image) file.

With the Primary/Secondary flash option you can test a new image in your system without having to replace a previously existing image. You can also use the image options for troubleshooting. For example, you can copy a problem image into Secondary flash for later analysis and place another, proven image in Primary flash to run your system. The switch can use only one image at a time.

The following tasks involve primary/secondary flash options:

- Displaying the current flash image data and determining which OS versions are available
- OS downloads
- Local OS replacement, and removal (erasing)
- System booting

Displaying the Current Flash Image Data

Use the commands in this section to:

- Determine whether there are flash images in both primary and secondary flash
- Determine whether the images in primary and secondary flash are the same
- Identify which OS version is currently running

Viewing the Currently Active Flash Image Version. This command identifies the software version on which the switch is currently running, and whether the active version was booted from the primary or secondary flash image.

Syntax: `show version`

For example, if the switch is using an OS version of E.05.01 stored in Primary flash, **show version** produces the following:

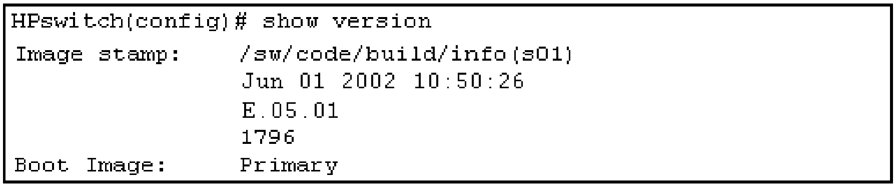

```
HPswitch(config)# show version
 Image stamp:    /sw/code/build/info(s01)
                 Jun 01 2002 10:50:26
                 E.05.01
                 1796
Boot Image:      Primary
```

Figure 5-6. Example Showing the Identity of the Current Flash Image

Determining Whether the Flash Images Are Different Versions. If the flash image sizes in primary and secondary are the same, then in almost every case, the primary and secondary images are identical. This command provides a comparison of flash image sizes, plus the boot ROM version and from which flash image the switch booted. For example, in the following case, the images are different versions of the OS software (flash image) and the switch is running on the version stored in the secondary flash image:

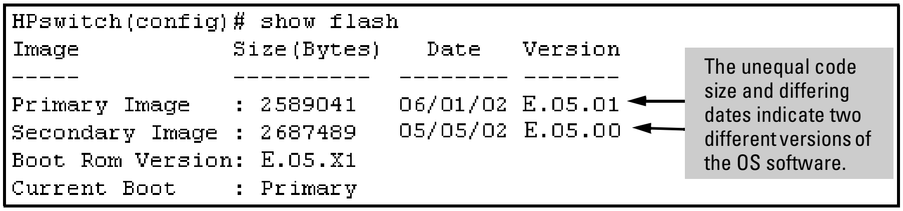

```
HPswitch(config)# show flash
Image           Size(Bytes)   Date     Version
-----           ----------  -------- -------
Primary Image   : 2589041   06/01/02 E.05.01
Secondary Image : 2687489   05/05/02 E.05.00
Boot Rom Version: E.05.X1
Current Boot    : Primary
```

The unequal code size and differing dates indicate two different versions of the OS software.

Figure 5-7. Example Showing Different Flash Image Versions

Determining Which Flash Image Versions Are Installed. The **show version** command displays which software version the switch is currently running and whether that version booted from primary or secondary flash. Thus, if the switch booted from primary flash, you will see the version number of the OS image stored in primary flash, and if the switch booted from secondary flash, you will see the version number of the OS version stored in secondary flash. Thus, by using **show version**, then rebooting the switch from the opposite flash image and using **show version** again, you can determine the version of the OS image in both flash sources. For example:

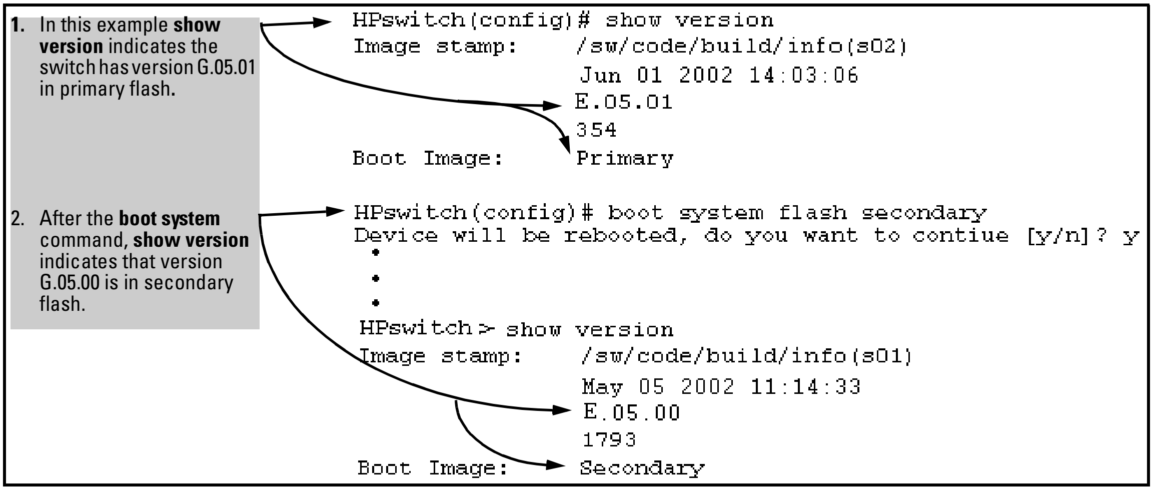

1. In this example **show version** indicates the switch has version G.05.01 in primary flash.

```
HPswitch(config)# show version
Image stamp:    /sw/code/build/info(s02)
                 Jun 01 2002 14:03:06
                E.05.01
                354
Boot Image:     Primary
```

2. After the **boot system** command, **show version** indicates that version G.05.00 is in secondary flash.

```
HPswitch(config)# boot system flash secondary
Device will be rebooted, do you want to contiue [y/n]? y
 •
 •
 •
HPswitch > show version
Image stamp:    /sw/code/build/info(s01)
                May 05 2002 11:14:33
                E.05.00
                1793
Boot Image:     Secondary
```

Figure 5-8. Determining the OS Version in Primary and Secondary Flash

OS Downloads

The following table shows the switch's options for downloading an OS to flash and booting the switch from flash

Table 5-1. Primary/Secondary Memory Access

Action	Menu	CLI	Web Browser	SNMP
Download to Primary	Yes	Yes	Yes	Yes
Download to Secondary	No	Yes	No	Yes
Boot from Primary	Yes	Yes	Yes	Yes
Boot from Secondary	No	Yes	No	Yes

The different OS download options involve different **copy** commands, plus **xmodem**, and **tftp**. These topics are covered in appendix A, "File Transfers".

Download Interruptions. In most cases, if a power failure or other cause interrupts a flash image download, the switch reboots with the image previously stored in primary flash. In the unlikely event that the primary image is corrupted, as a result of an interruption, the switch will reboot from secondary flash and you can either copy the secondary image into primary or download another image to primary from an external source. See Appendix A, "File Transfers".

Local OS Replacement and Removal

This section describes commands for erasing an OS (flash image) and copying an existing OS between primary and secondary flash.

Note

It is not necessary to erase the content of a flash location before downloading another OS file. The process automatically overwrites the previous file with the new file. If you want to remove an unwanted OS version from flash, HP recommends that you do so by overwriting it with the same OS version that you are using to operate the switch, or with another acceptable OS version. To copy an OS image file between the primary and secondary flash locations, see "Copying an OS Image from One Flash Location to Another" , below.

The local commands described here are for flash image management within the switch. To download an OS image file from an external source, see Appendix A, "File Transfers".

Copying an OS Image from One Flash Location to Another. When you copy the flash image from primary to secondary or the reverse, the switch overwrites the file in the destination location with a copy of the file from the source location. This means you *do not* have to erase the current image at the destination location before copying in a new image.

Caution

Verify that there is an acceptable OS image in the source flash location from which you are going to copy. Use the **show flash** command or, if necessary, the procedure under "Determining Which Flash Image Versions Are Installed" on page 5-13 to verify an acceptable OS image. Attempting to copy from a source image location that has a corrupted flash image overwrites the image in the destination flash location. In this case, the switch will not have a valid flash image in either flash location, but will continue running on a temporary flash image in RAM. *Do not reboot the switch.* Instead, immediately download another valid flash image to primary or secondary flash. Otherwise, if the switch is rebooted without an OS image in either primary or secondary flash, the temporary flash image in RAM will be cleared and the switch will go down. To recover, see "Restoring a Flash Image" on page C-35 (in the Troubleshooting chapter).

Syntax: copy flash flash <*destination flash*>

where: *destination flash* = **primary** or **secondary**:

For example, to copy the image in secondary flash to primary flash:

1. Verify that there is a valid flash image in the secondary flash location. The following figure indicates that an OS image is present in secondary flash. (If you are unsure whether the image is secondary flash is valid, try booting from it before you proceed, by using **boot system flash secondary**.)

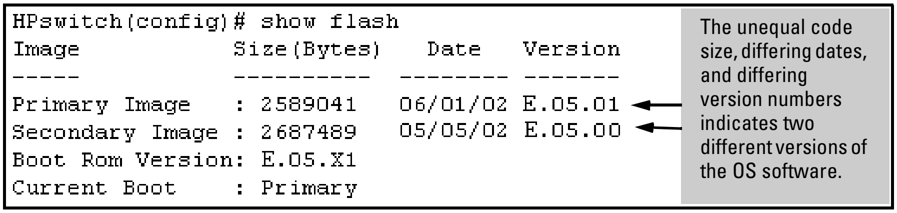

```
HPswitch(config)# show flash
Image             Size(Bytes)   Date     Version
-----             ----------  -------- -------
Primary Image   : 2589041     06/01/02 E.05.01
Secondary Image : 2687489     05/05/02 E.05.00
Boot Rom Version: E.05.X1
Current Boot    : Primary
```

The unequal code size, differing dates, and differing version numbers indicates two different versions of the OS software.

Figure 5-9. Example Indicating Two Different OS Versions in Primary and Secondary Flash

Execute the copy command as follows:

```
HPswitch(config)# copy flash flash primary
```

Erasing the Contents of Primary or Secondary Flash. This command deletes the OS image file from the specified flash location.

Caution--No Undo!

Before using this command in one flash image location (primary or secondary), ensure that you have a valid OS file in the other flash image location (secondary or primary). If the switch has only one flash image loaded (in either primary or secondary flash) and you erase that image, then the switch does not have an OS stored in flash. In this case, if you do not reboot or power cycle the switch, you can recover by using xmodem or tftp to download another OS.

Syntax: erase flash < primary | secondary >

For example, to erase the OS in primary flash, do the following:

1. First verify that a usable flash image exists in secondary flash. The most reliable way to ensure this is to reboot the switch from the flash image you want to retain. For example, if you are planning to erase the primary image, then first reboot from the secondary image to verify that the secondary image is present and acceptable for your system:

   ```
   HPswitch# boot system flash secondary
   ```

2. Then erase the OS in the selected flash (in this case, primary):

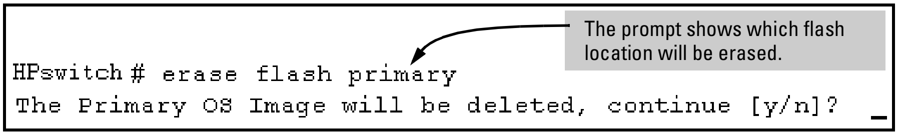

```
HPswitch# erase flash primary
 The Primary OS Image will be deleted, continue [y/n]? _
```

The prompt shows which flash location will be erased.

Figure 5-10. Example of Erase Flash Prompt

3. Type **y** at the prompt to complete the flash erase.
4. Use **show flash** to verify erasure of the selected OS flash image

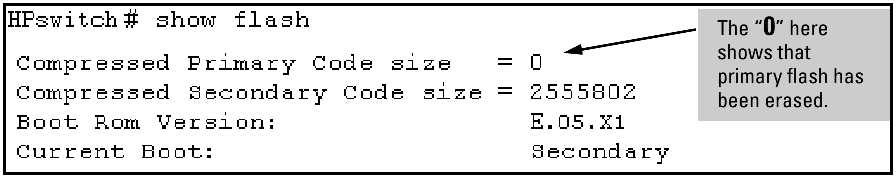

```
HPswitch# show flash

 Compressed Primary Code size   = 0
 Compressed Secondary Code size = 2555802
 Boot Rom Version:                E.05.X1
 Current Boot:                    Secondary
```

The "**0**" here shows that primary flash has been erased.

Figure 5-11. Example of Show Flash Listing After Erasing Primary Flash

Rebooting the Switch

The switch offers reboot options through the **boot** and **reload** commands, plus the options inherent in a dual-flash image system. Generally, using **boot** provides more comprehensive self-testing; using **reload** gives you a faster reboot time.

Table 5-2. Comparing the Boot and Reload Commands

Actions	Included In Boot?	Included In Reload	Note
Save all configuration changes since the last boot or reload	Optional, with prompt	Yes, automatic	Config changes saved to the startup-config file
Perform all system self-tests	Yes	No	Reload provides a faster system reboot.
Choice of primary or secondary	Yes	No—Uses the current flash image.	

Booting from Primary Flash. This command always boots the switch from primary flash, and executes the complete set of subsystem self-tests.

Syntax: boot

For example, to boot the switch from primary flash with pending configuration changes in the running-config file:

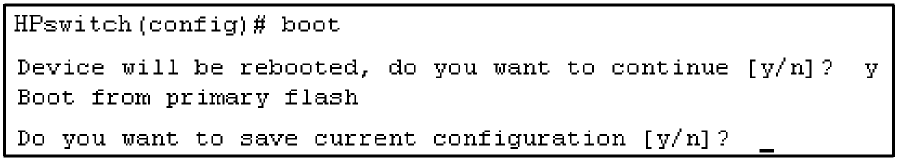

```
HPswitch(config)# boot

Device will be rebooted, do you want to continue [y/n]?  y
Boot from primary flash

Do you want to save current configuration [y/n]? _
```

Figure 5-12. Example of Boot Command (Default Primary Flash)

In the above example, typing either a **y** or **n** at the second prompt initiates the reboot operation. Also, if there are no pending configuration changes in the running-config file, then the reboot commences without the pause to display `Boot from primary flash`.

Booting from a Specified Flash. This version of the boot command gives you the option of specifying whether to reboot from primary or secondary flash, and is the required command for rebooting from secondary flash. This option also executes the complete set of subsystem self-tests.

Syntax: boot system flash < primary | secondary >

For example, to reboot the switch from secondary flash when there are no pending configuration changes in the running-config file:

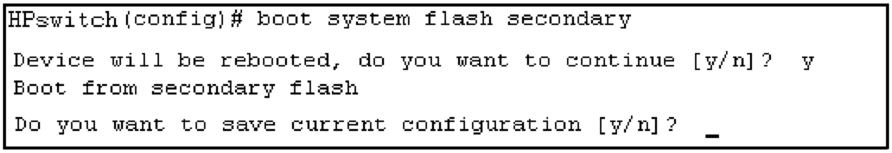

```
HPswitch(config)# boot system flash secondary

Device will be rebooted, do you want to continue [y/n]?  y
Boot from secondary flash

Do you want to save current configuration [y/n]? _
```

Figure 5-13. Example of Boot Command with Primary/Secondary Flash Option

In the above example, typing either a **y** or **n** at the second prompt initiates the reboot operation. Also, if there are no pending configuration changes in the running-config file, then the reboot commences without the pause to display `Boot from secondary flash`.

Booting from the Current OS Version. **Reload** reboots the switch from the flash image on which the switch is currently running, and saves to the startup-config file any configuration changes currently in the running-config file. Because **reload** bypasses some subsystem self-tests, the switch reboots faster than if you use either of the **boot** command options.

Syntax: reload

For example, if you change the number of VLANs the switch supports, you must reboot the switch in order to implement the change. Reload automatically saves your configuration changes and reboots the switch from the same OS you have been using:

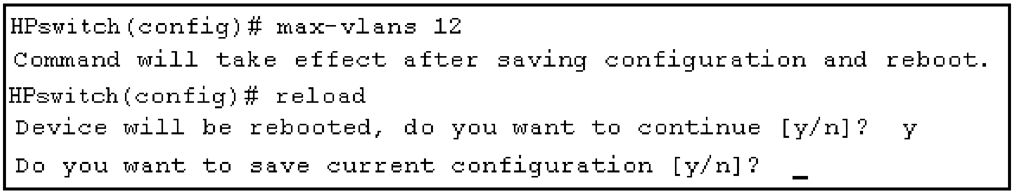

```
HPswitch(config)# max-vlans 12
Command will take effect after saving configuration and reboot.
HPswitch(config)# reload
Device will be rebooted, do you want to continue [y/n]?  y
Do you want to save current configuration [y/n]? _
```

Figure 5-14. Using Reload with Pending Configuration Changes

Operating Notes

Default Boot Source. The switch reboots from primary flash by default unless you specify the secondary flash.

Boot Attempts from an Empty Flash Location. In this case, the switch aborts the attempt and displays

```
Image does not exist
Operation aborted.
```

Interaction of Primary and Secondary Flash Images with the Current Configuration. The switch has one startup-config file (page 5-2), which it always uses for reboots, regardless of whether the reboot is from primary or secondary flash. Also, for rebooting purposes, it is not necessary for the OS and the startup-config file to support identical software features. For example, suppose you have just downloaded an OS upgrade that includes new features that are not supported in the OS you used to create the current startup-config file. In this case, the OS simply assigns factory-default values to the parameters controlling the new features. Similarly, If you create a startup-config file while using a version “Y” of the OS, and then reboot the switch with an earlier OS version “X” that does not include all of the features found in “Y”, the OS simply ignores the parameters for any features that it does not support.

Interface Access, System Information, and Friendly Port Names

Contents

Overview

This chapter describes how to:

- View and modify the configuration for switch interface access
- Use the CLI **kill** command to terminate a remote session
- View and modify switch system information

For help on how to actually use the interfaces built into the switch, refer to:

- Chapter 2, “Using the Menu Interface”
- Chapter 3, “Using the Command Line Interface (CLI)”
- Chapter 4, Using the HP Web Browser Interface”

Why Configure Interface Access and System Information? The interface access features in the switch operate properly by default. However, you can modify or disable access features to suit your particular needs. Similarly, you can choose to leave the system information parameters at their default settings. However, modifying these parameters can help you to more easily distinguish one device from another in your network.

Interface Access: Console/Serial Link, Web, and Inbound Telnet

Interface Access Features

Feature	Default	Menu	CLI	Web
Inactivity Time	0 Minutes (disabled)	page 6-4	page 6-6	—
Inbound Telnet Access	Enabled	page 6-4	page 6-5	—
Outbound Telnet Access	n/a	—	page 6-6	—
Web Browser Interface Access	Enabled	page 6-4	page 6-6	—
Terminal type	VT-100	—	page 6-6	—
Event Log event types to list (Displayed Events)	All	—	page 6-6	—
Baud Rate	Speed Sense	—	page 6-6	—
Flow Control	XON/XOFF	—	page 6-6	—

In most cases, the default configuration is acceptable for standard operation.

Note

Basic switch security is through passwords. You can gain additional security using IP authorized managers. However if unauthorized access to the switch through in-band means (Telnet or the web browser interface), then you can disallow in-band access (as described in this section) and install the switch in a locked environment.

Menu: Modifying the Interface Access

The menu interface enables you to modify these parameters:

- Inactivity Timeout
- Inbound Telnet Enabled
- Web Agent Enabled

To Access the Interface Access Parameters:

1. From the Main Menu, Select...

   2. Switch Configuration...

   1. System Information

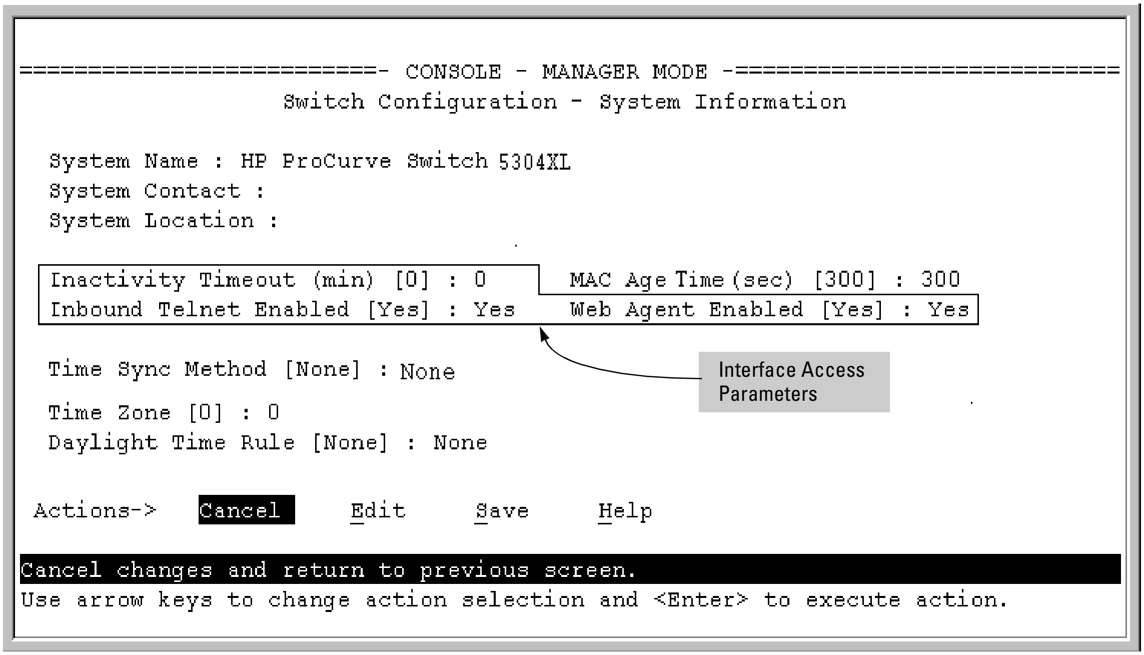

```
========================- CONSOLE - MANAGER MODE -===========================
                   Switch Configuration - System Information

  System Name : HP ProCurve Switch 5304XL
  System Contact :
  System Location :

  Inactivity Timeout (min) [0] : 0      MAC Age Time(sec) [300] : 300
  Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes

  Time Sync Method [None] : None

  Time Zone [0] : 0
  Daylight Time Rule [None] : None

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Interface Access Parameters

Figure 6-1. The Default Interface Access Parameters Available in the Menu Interface

2. Press [E] (for Edit). The cursor moves to the **System Name** field.
3. Use the arrow keys ([↓], [↑], [←], [→]) to move to the parameters you want to change.

   Refer to the online help provided with this screen for further information on configuration options for these features.
4. When you have finished making changes to the above parameters, press [Enter], then press [S] (for **Save**).

CLI: Modifying the Interface Access

Interface Access Commands Used in This Section

show console	below
[no] telnet-server	below
[no] web-management	page 6-6
console	page 6-6

Listing the Current Console/Serial Link Configuration. This command lists the current interface access parameter settings.

Syntax: show console

This example shows the switch's default console/serial configuration.

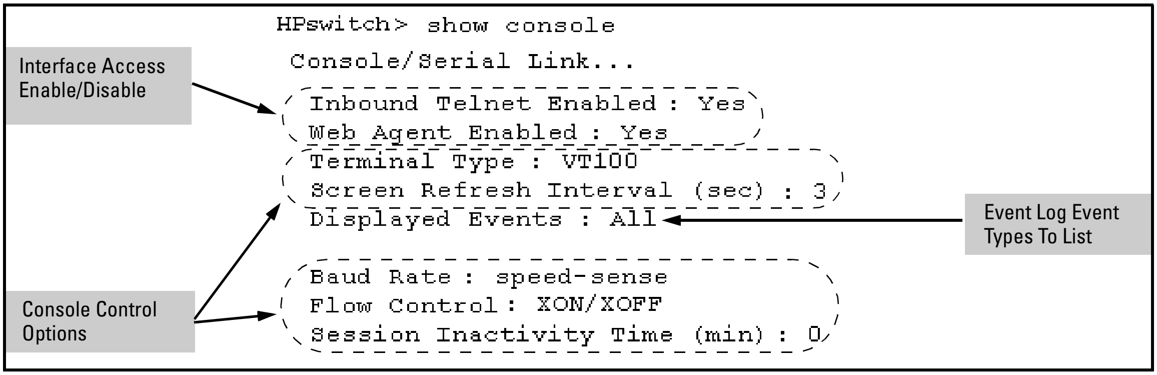

```
HPswitch> show console
 Console/Serial Link...
  Inbound Telnet Enabled : Yes
  Web Agent Enabled : Yes
  Terminal Type : VT100
  Screen Refresh Interval (sec) : 3
  Displayed Events : All

  Baud Rate : speed-sense
  Flow Control : XON/XOFF
  Session Inactivity Time (min) : 0
```

Interface Access Enable/Disable
Console Control Options
Event Log Event Types To List

Figure 6-2. Listing of Show Console Command

Reconfigure Inbound Telnet Access. In the default configuration, inbound Telnet access is enabled.

Syntax: [no] telnet-server

To disable inbound Telnet access:

```
HPswitch(config)# no telnet-server
```

To re-enable inbound Telnet access:

```
HPswitch(config)# telnet-server
```

Outbound Telnet to Another Device. This feature operates independently of the telnet-server status and enables you to Telnet to another device that has an IP address.

Syntax: telnet < *ip-address* >

For example:

```
HPswitch # telnet 10.28.27.204
```

Reconfigure Web Browser Access. In the default configuration, web browser access is enabled.

Syntax:[no] web-management

To disable web browser access:

```
HPswitch(config)# no web-management
```

To re-enable web browser access:

```
HPswitch(config)# web-management
```

Reconfigure the Console/Serial Link Settings. You can reconfigure one or more console parameters with one console command.

Syntax:

```
console
   [terminal <vt100 | ansi>]
   [screen-refresh <1 | 3 | 5 | 10 | 20 | 30 | 45 | 60>]
   [baud-rate
        <speed-sense | 1200 | 2400 | 4800 | 9600 | 19200 |38400 | 57600>]
   [flow-control <xon/xoff | none>]
   [inactivity-timer <0  1  5  10  15  20  30  60  120>]
   [events <none | all | non-info | critical | debug]
```

Note

If you change the Baud Rate or Flow Control settings for the switch, you should make the corresponding changes in your console access device. Otherwise, you may lose connectivity between the switch and your terminal emulator due to differences between the terminal and switch settings for these two parameters.

All console parameter changes except **events** require that you save the configuration with **write memory** and then execute **boot** before the new console configuration will take effect.

For example, to use one command to configure the switch with the following:

- VT100 operation
- 19,200 baud
- No flow control
- 10-minute inactivity time
- Critical log events

you would use the following command sequence:

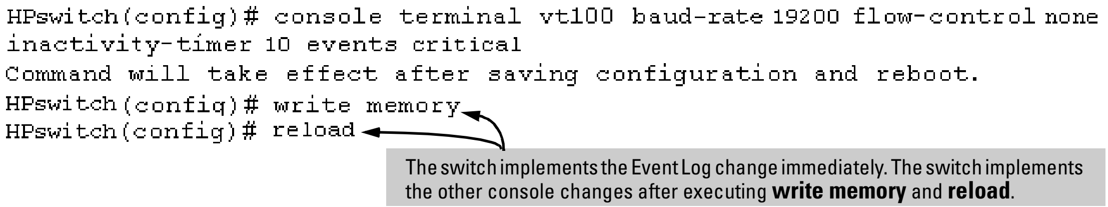

```
HPswitch(config)# console terminal vt100 baud-rate 19200 flow-control none
inactivity-timer 10 events critical
Command will take effect after saving configuration and reboot.
HPswitch(config)# write memory
HPswitch(config)# reload
```

The switch implements the Event Log change immediately. The switch implements the other console changes after executing **write memory** and **reload**.

Figure 6-3. Example of Executing the Console Command with Multiple Parameters

You can also execute a series of console commands and then save the configuration and boot the switch. For example:

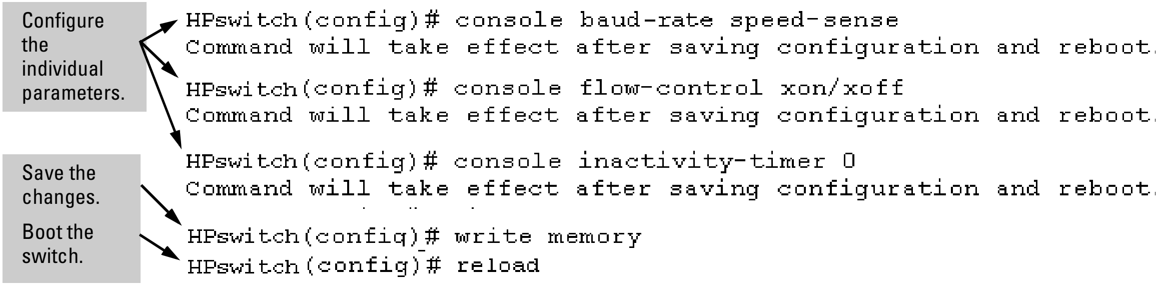

Configure the individual parameters.

```
HPswitch(config)# console baud-rate speed-sense
Command will take effect after saving configuration and reboot.
HPswitch(config)# console flow-control xon/xoff
Command will take effect after saving configuration and reboot.
HPswitch(config)# console inactivity-timer 0
Command will take effect after saving configuration and reboot.
HPswitch(config)# write memory
HPswitch(config)# reload
```

Save the changes.

Boot the switch.

Figure 6-4. Example of Executing a Series of Console Commands

Denying Interface Access by Terminating Remote Management Sessions

The switch supports up to four management sessions. You can use **show ip ssh** to list the current management sessions, and **kill** to terminate a currently running remote session. (**Kill** does not terminate a Console session on the serial port, either through a direct connection or via a modem.)

Syntax: kill [<*session-number*>]

For example, if you are using the switch's serial port for a console session and want to terminate a currently active Telnet session, you would do the following:

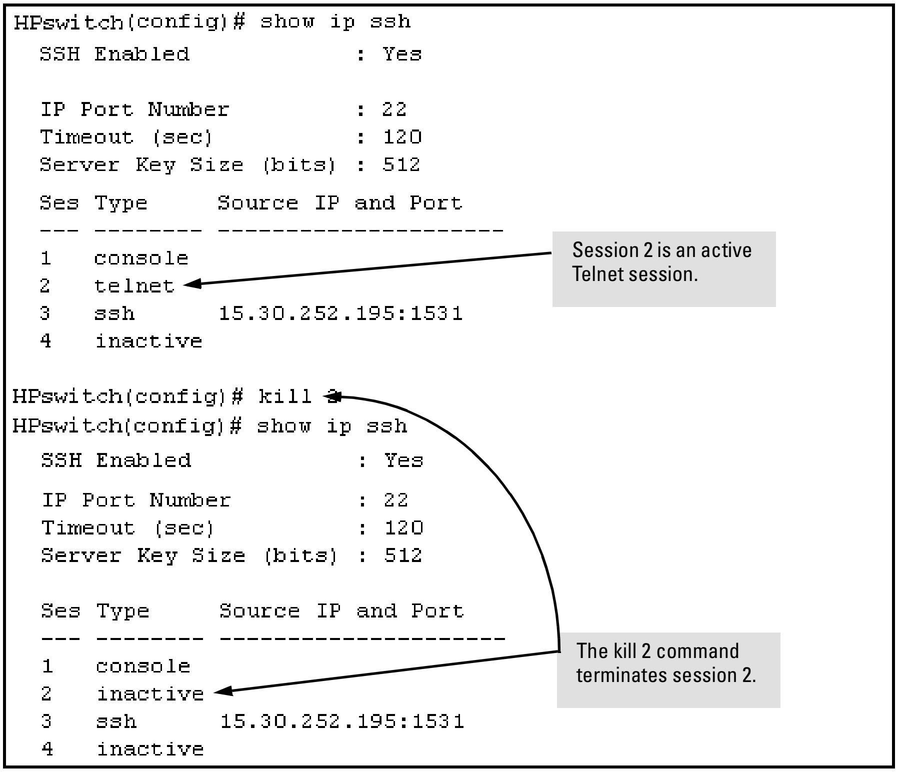

```
HPswitch(config)# show ip ssh
  SSH Enabled            : Yes

  IP Port Number         : 22
  Timeout (sec)          : 120
  Server Key Size (bits) : 512

  Ses Type     Source IP and Port
  --- -------- ---------------------
  1   console
  2   telnet
  3   ssh      15.30.252.195:1531
  4   inactive

HPswitch(config)# kill 2
HPswitch(config)# show ip ssh
  SSH Enabled            : Yes

  IP Port Number         : 22
  Timeout (sec)          : 120
  Server Key Size (bits) : 512

  Ses Type     Source IP and Port
  --- -------- ---------------------
  1   console
  2   inactive
  3   ssh      15.30.252.195:1531
  4   inactive
```

Session 2 is an active Telnet session.

The kill 2 command terminates session 2.

Figure 6-5. Example of Using the "Kill" Command To Terminate a Remote Session

System Information

System Information Features

Feature	Default	Menu	CLI	Web
System Name	*switch product name*	page 6-10	page 6-12	page 6-14
System Contact	n/a	page 6-10	page 6-12	page 6-14
System Location	n/a	page 6-10	page 6-12	page 6-14
MAC Age Time	300 seconds	page 6-10	page 6-13	—
Time Sync Method	None	See Chapter 8, "Time Protocols".		
Time Zone	0	page 6-10	page 6-13	—
Daylight Time Rule	None	page 6-10	page 6-13	—
Time	January 1, 1990 at 00:00:00 at last power reset	—	page 6-13	—

Configuring system information is optional, but recommended.

System Name: Using a unique name helps you to identify individual devices where you are using an SNMP network management tool such as HP TopTools for Hubs & Switches.

System Contact and Location: This information is helpful for identifying the person administratively responsible for the switch and for identifying the locations of individual switches.

MAC Age Time: The number of seconds a MAC address the switch has learned remains in the switch's address table before being aged out (deleted). Aging out occurs when there has been no traffic from the device belonging to that MAC address for the configured interval.

Time Sync Method: Selects the method (TimeP or SNTP) the switch will use for time synchronization. For more on this topic, refer to Chapter 8, "Time Protocols".

Time Zone: The number of minutes your time zone location is to the West (+) or East (-) of Coordinated Universal Time (formerly GMT). The default **0** means no time zone is configured. For example, the time zone for Berlin, Germany is + 60 (minutes) and the time zone for Vancouver, Canada is - 480 (minutes).

Daylight Time Rule: Specifies the daylight savings time rule to apply for your location. The default is **None**. (For more on this topic, see appendix D, "Daylight Savings Time on HP ProCurve Switches.)

Time: Used in the CLI to specify the time of day, the date, and other system parameters.

Menu: Viewing and Configuring System Information

To access the system information parameters:

1. From the Main Menu, Select...

   2. Switch Configuration...

   1. System Information

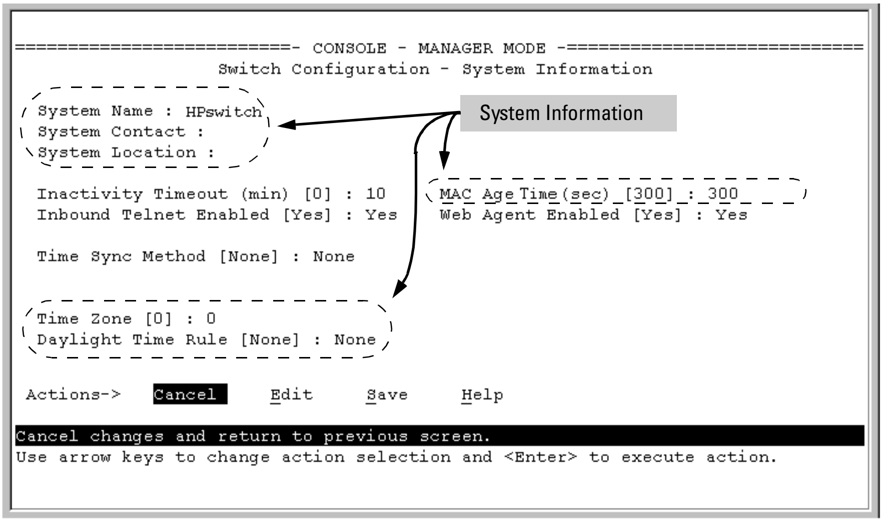

```
=========================- CONSOLE - MANAGER MODE -============================
                   Switch Configuration - System Information

  System Name : HPswitch
  System Contact :
  System Location :

  Inactivity Timeout (min) [0] : 10       MAC Age Time (sec) [300] : 300
  Inbound Telnet Enabled [Yes] : Yes      Web Agent Enabled [Yes] : Yes

  Time Sync Method [None] : None

  Time Zone [0] : 0
  Daylight Time Rule [None] : None

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 6-6. The System Information Configuration Screen (Default Values)

Note

To help simplify administration, it is recommended that you configure **System Name** to a character string that is meaningful within your system.

2. Press E (for Edit). The cursor moves to the **System Name** field.
3. Refer to the online help provided with this screen for further information on configuration options for these features.
4. When you have finished making changes to the above parameters, press Enter, then press S (for **Save**) and return to the Main Menu.

CLI: Viewing and Configuring System Information

System Information Commands Used in This Section

show system-information	below
hostname	below
snmp-server [contact] [location]	below
mac-age-time	page 6-13
time	
timezone	page 6-13
daylight-time-rule	page 6-13
date time	page 6-13

Listing the Current System Information. This command lists the current system information settings.

Syntax: show system-information

This example shows the switch's default console configuration.

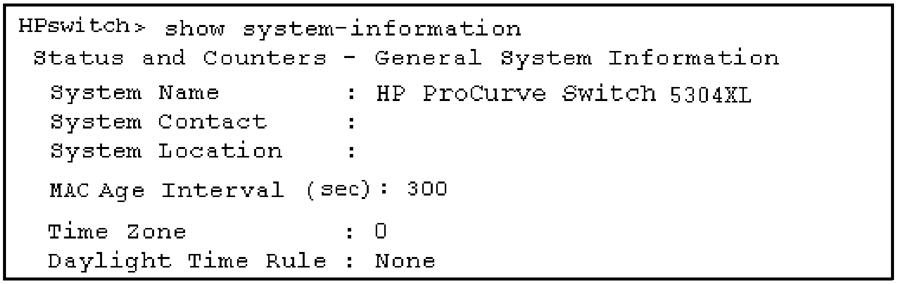

```
HPswitch> show system-information
 Status and Counters - General System Information

  System Name         : HP ProCurve Switch 5304XL
  System Contact      :
  System Location     :

  MAC Age Interval (sec): 300

  Time Zone           : 0
  Daylight Time Rule  : None
```

Figure 6-7. Example of CLI System Information Listing

Configure a System Name, Contact, and Location for the Switch. To help distinguish one switch from another, configure a plain-language identity for the switch.

Syntax: hostname <*name-string*>
snmp-server [contact <*system contact*>] [location <*system location*>]

Both fields allow up to 48 characters. *Blank spaces* are not allowed in the variables for these commands.

For example, to name the switch “Blue” with “Next-4474” as the system contact, and “North-Data-Room” as the location:

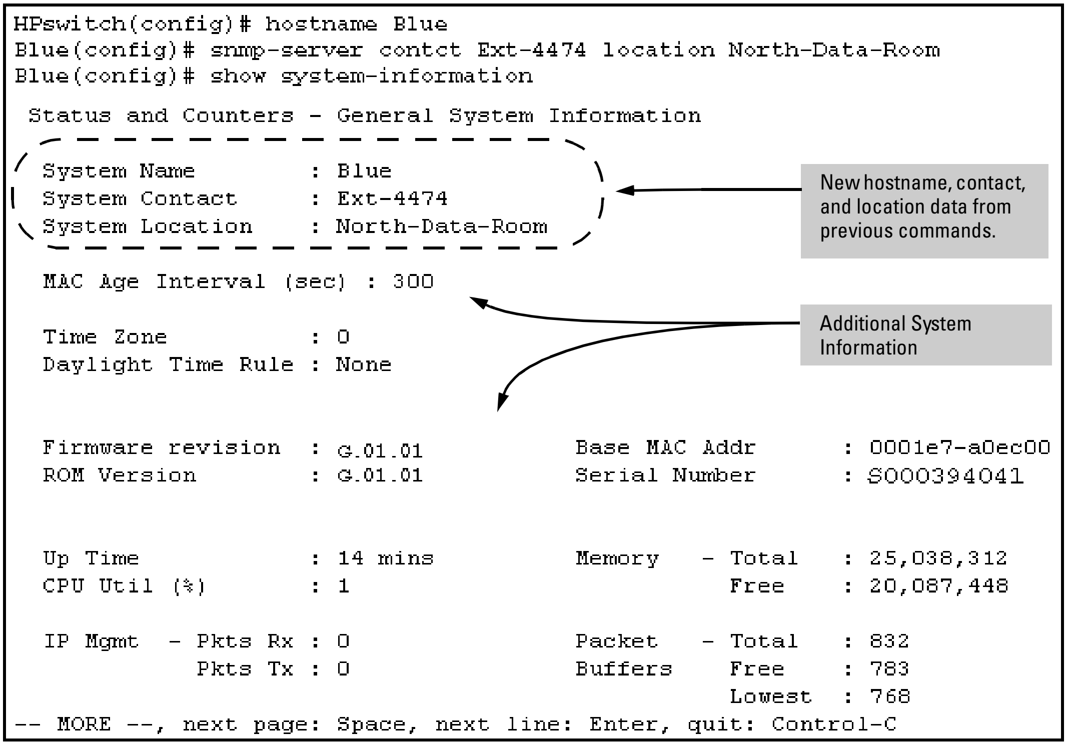

```
HPswitch(config)# hostname Blue
Blue(config)# snmp-server contct Ext-4474 location North-Data-Room
Blue(config)# show system-information

 Status and Counters - General System Information

  System Name         : Blue
  System Contact      : Ext-4474
  System Location     : North-Data-Room

  MAC Age Interval (sec) : 300

  Time Zone           : 0
  Daylight Time Rule : None


  Firmware revision   : G.01.01          Base MAC Addr      : 0001e7-a0ec00
  ROM Version         : G.01.01          Serial Number      : S000394041


  Up Time             : 14 mins          Memory   - Total   : 25,038,312
  CPU Util (%)        : 1                           Free    : 20,087,448

  IP Mgmt  - Pkts Rx : 0                 Packet   - Total   : 832
             Pkts Tx : 0                 Buffers    Free    : 783
                                                    Lowest  : 768
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

New hostname, contact, and location data from previous commands.

Additional System Information

Figure 6-8. System Information Listing After Executing the Preceding Commands

Reconfigure the Age Time for Learned MAC Addresses. This command corresponds to the MAC Age Interval in the menu interface, and is expressed in seconds.

Syntax: mac-age-time <*10 . . 1000000*> (seconds)

For example, to configure the age time to seven minutes:

```
HPswitch(config)# mac-age-time 420
```

Configure the Time Zone and Daylight Time Rule. These commands:

- Set the time zone you want to use
- Define the daylight time rule for keeping the correct time when daylight-saving-time shifts occur.

Syntax: time timezone <-720 . . 840>
time daylight-time-rule <none | alaska | continental-us-and-canada | middle-europe-and-portugal | southern-hemisphere | western-europe | user-defined>

East of the 0 meridian, the sign is "+". West of the 0 meridian, the sign is "-".

For example, the time zone setting for Berlin, Germany is +60 (zone +1, or 60 minutes), and the time zone setting for Vancouver, Canada is -480 (zone -8, or -480 minutes). To configure the time zone and daylight time rule for Vancouver, Canada:

```
HPswitch(config)# time timezone -480 daylight-time-rule
   continental-us-and-canada
```

Configure the Time and Date. The switch uses the time command to configure both the time of day and the date. Also, executing time without parameters lists the switch's time of day and date. Note that the CLI uses a 24-hour clock scheme; that is, hour (*hh*) values from 1 p.m. to midnight are input as 13 - 24, respectively.

Syntax: time [*hh:mm*[:*ss*]] [*mm/dd/* [*yy*]*yy*]

For example, to set the switch to 9:45 a.m. on November 17, 2002:

```
HPswitch(config)# time 9:45 11/17/02
```

Note

Executing **reload** or **boot** resets the time and date to their default startup values.

Web: Configuring System Parameters

In the web browser interface, you can enter the following system information:

- System Name
- System Location
- System Contact

For access to the MAC Age Interval and the Time parameters, use the menu interface or the CLI.

Configure System Parameters in the Web Browser Interface.

1. Click on the **Configuration** tab.
2. Click on [System Info].
3. Enter the data you want in the displayed fields.
4. Implement your new data by clicking on [Apply Changes].

To access the web-based help provided for the switch, click on [?] in the web browser screen.

Using Friendly (Optional) Port Names

Feature	Default	Menu	CLI	Web
Configure Friendly Port Names	Standard Port Numbering	n/a	page 16	n/a
Display Friendly Port Names	n/a	n/a	page 18	n/a

This feature enables you to assign alphanumeric port names of your choosing to augment automatically assigned numeric port names. This means you can configure meaningful port names to make it easier to identify the source of information listed by some **Show** commands. (Note that this feature *augments* port numbering, but *does not replace* it.)

Configuring and Operating Rules for Friendly Port Names

- At either the global or context configuration level you can assign a unique name to any port on the switch. You can also assign the same name to multiple ports.
- The friendly port names you configure appear in the output of the **show name [*port-list*]**, **show config**, and **show interface <*port-number*>** commands. They do not appear in the output of other show commands or in Menu interface screens. (See "Displaying Friendly Port Names with Other Port Data" on page 6-18.)
- Friendly port names are not a substitute for port numbers in CLI commands or Menu displays.
- Trunking ports together does not affect friendly naming for the individual ports. (If you want the same name for all ports in a trunk, you must individually assign the name to each port.)
- A friendly port name can have up to 64 contiguous alphanumeric characters.
- Blank spaces within friendly port names are not allowed, and if used, cause an **invalid input** error. (The switch interprets a blank space as a name terminator.)
- In a port listing, **not assigned** indicates that the port does not have a name assignment other than its fixed port number.

- To retain friendly port names across reboots, you must save the current running-configuration to the startup-config file after entering the friendly port names. (In the CLI, use the **write memory** command.)

Configuring Friendly Port Names

Syntax:interface [e] <*port-list*> name <*port-name-string*>Assigns a port name to *port-list*.

no interface [e] <*port-list*> nameDeletes the port name from *port-list*.

Configuring a Single Port Name. Suppose that you have connected port A3 on the switch to Bill Smith's workstation, and want to assign Bill's name and workstation IP address (10.25.101.73) as a port name for port A3:

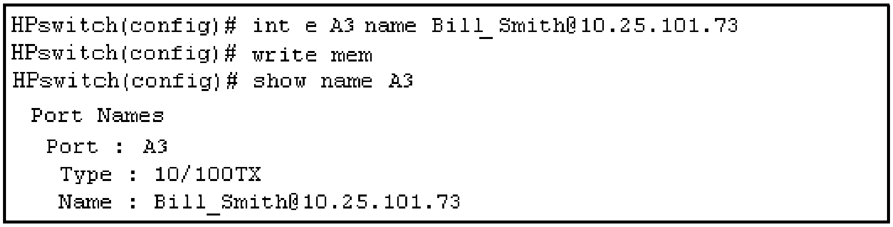

```
HPswitch(config)# int e A3 name Bill_Smith@10.25.101.73
HPswitch(config)# write mem
HPswitch(config)# show name A3
 Port Names
  Port : A3
   Type : 10/100TX
   Name : Bill_Smith@10.25.101.73
```

Figure 6-9. Example of Configuring a Friendly Port Name

Configuring the Same Name for Multiple Ports. Suppose that you want to use ports A5 through A8 as a trunked link to a server used by a drafting group. In this case you might configure ports A5 through A8 with the name "Draft-Server:Trunk".

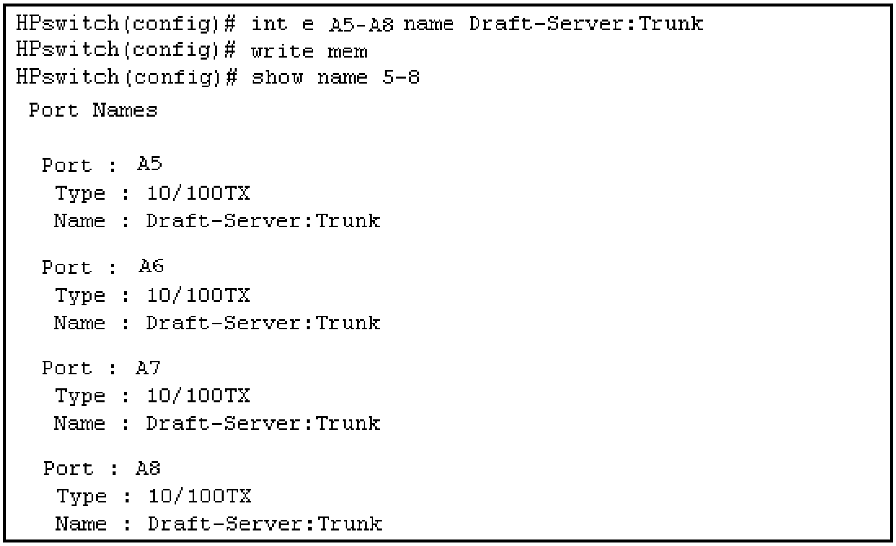

```
HPswitch(config)# int e A5-A8 name Draft-Server:Trunk
HPswitch(config)# write mem
HPswitch(config)# show name 5-8
 Port Names

  Port :  A5
   Type : 10/100TX
   Name : Draft-Server:Trunk

  Port :  A6
   Type : 10/100TX
   Name : Draft-Server:Trunk

  Port :  A7
   Type : 10/100TX
   Name : Draft-Server:Trunk

  Port : A8
   Type : 10/100TX
   Name : Draft-Server:Trunk
```

Figure 6-10. Example of Configuring One Friendly Port Name on Multiple Ports

Displaying Friendly Port Names with Other Port Data

You can display friendly port name data in the following combinations:

- **show name**: Displays a listing of port numbers with their corresponding friendly port names and also quickly shows you which ports do not have friendly name assignments. (**show name** data comes from the running-config file.)
- **show interface <*port-number*>**: Displays the friendly port name, if any, along with the traffic statistics for that port. (The friendly port name data comes from the running-config file.)
- **show config**: Includes friendly port names in the per-port data of the resulting configuration listing. (**show config** data comes from the startup-config file.)

To List All Ports or Selected Ports with Their Friendly Port Names.

This command lists names assigned to a specific port.

Syntax: show name [*port-list*] Lists the friendly port name with its corresponding port number and port type. The **show name** alone command lists this data for all ports on the switch.

For example:

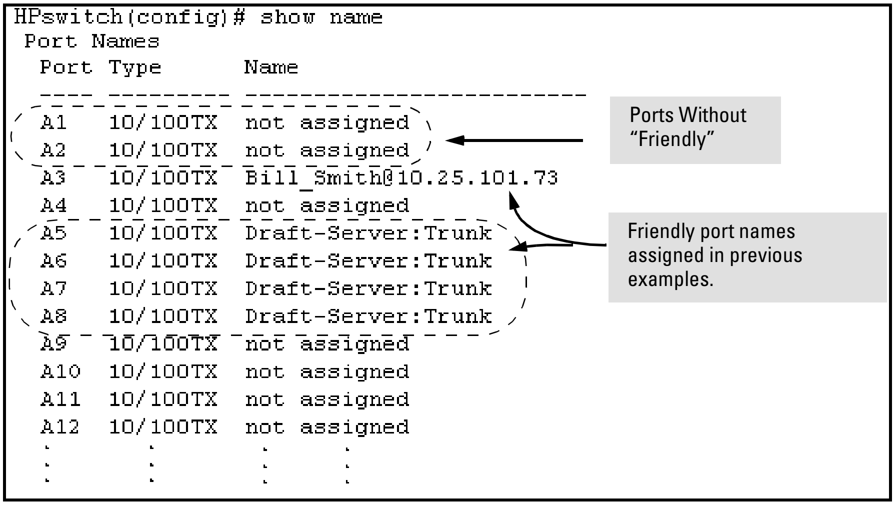

```
HPswitch(config)# show name
 Port Names
  Port Type      Name
  ---- --------- -------------------------
  A1   10/100TX  not assigned
  A2   10/100TX  not assigned
  A3   10/100TX  Bill_Smith@10.25.101.73
  A4   10/100TX  not assigned
  A5   10/100TX  Draft-Server:Trunk
  A6   10/100TX  Draft-Server:Trunk
  A7   10/100TX  Draft-Server:Trunk
  A8   10/100TX  Draft-Server:Trunk
  A9   10/100TX  not assigned
  A10  10/100TX  not assigned
  A11  10/100TX  not assigned
  A12  10/100TX  not assigned
  .    .          .     .
  .    .          .     .
  .    .          .     .
```

Ports Without "Friendly"

Friendly port names assigned in previous examples.

Figure 6-11. Example of Friendly Port Name Data for All Ports on the Switch

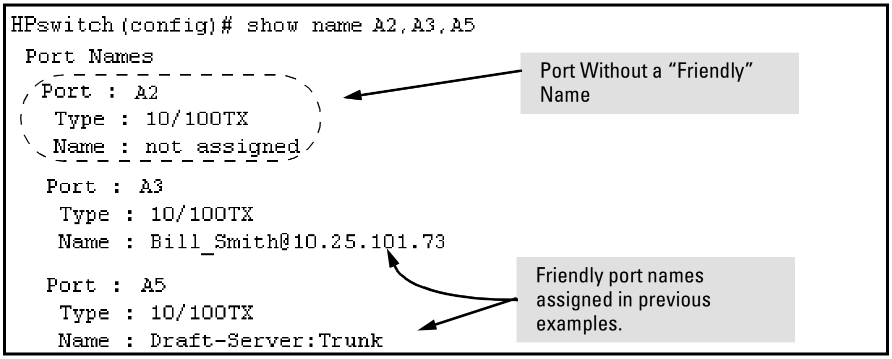

```
HPswitch(config)# show name A2,A3,A5
 Port Names
  Port : A2
   Type : 10/100TX
   Name : not assigned

   Port : A3
    Type : 10/100TX
    Name : Bill_Smith@10.25.101.73

   Port : A5
    Type : 10/100TX
    Name : Draft-Server:Trunk
```

Port Without a “Friendly” Name

Friendly port names assigned in previous examples.

Figure 6-12. Example of Friendly Port Name Data for Specific Ports on the Switch

Including Friendly Port Names in Per-Port Statistics Listings. A friendly port name configured to a port is automatically included when you display the port’s statistics output.

Syntax:	show interface <*port-number*>	Includes the friendly port name with the port’s traffic statistics listing.

For example, if you configure port A1 with the name “O’Connor_10.25.101.43”, the show interface output for this port appears similar to the following:

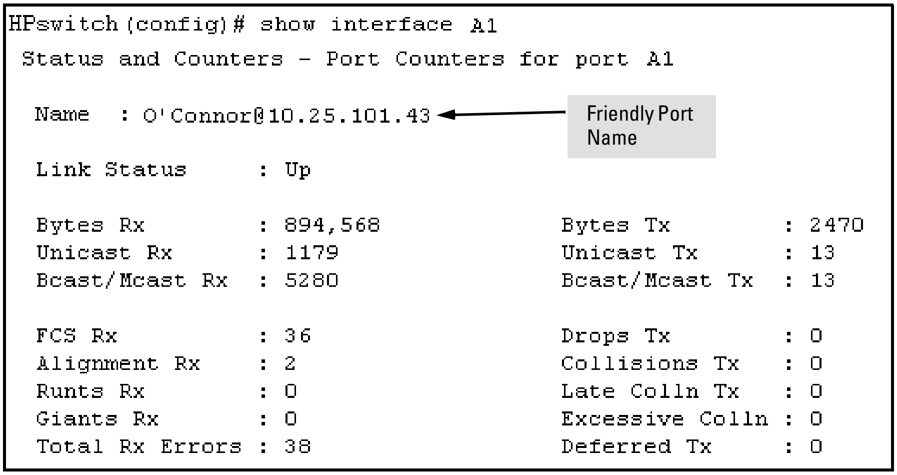

```
HPswitch(config)# show interface A1
 Status and Counters - Port Counters for port A1

  Name   : O'Connor@10.25.101.43

  Link Status     : Up

  Bytes Rx        : 894,568               Bytes Tx         : 2470
  Unicast Rx      : 1179                  Unicast Tx       : 13
  Bcast/Mcast Rx  : 5280                  Bcast/Mcast Tx   : 13

  FCS Rx          : 36                    Drops Tx         : 0
  Alignment Rx    : 2                     Collisions Tx    : 0
  Runts Rx        : 0                     Late Colln Tx    : 0
  Giants Rx       : 0                     Excessive Colln  : 0
  Total Rx Errors : 38                    Deferred Tx      : 0
```

Friendly Port Name

Figure 6-13. Example of a Friendly Port Name in a Per-Port Statistics Listing

For a given port, if a friendly port name does not exist in the running-config file, the `Name` line in the above command output appears as:

```
Name   :  not assigned
```

To Search the Configuration for Ports with Friendly Port Names.

This option tells you which friendly port names have been saved to the startup-config file. (**show config** does not include ports that have only default settings in the startup-config file.)

Syntax: show config — includes friendly port names in a listing of all interfaces (ports) configured with non-default settings. Excludes ports that have neither a friendly port name nor any other non-default configuration settings.

For example, if you configure port A1 with a friendly port name:

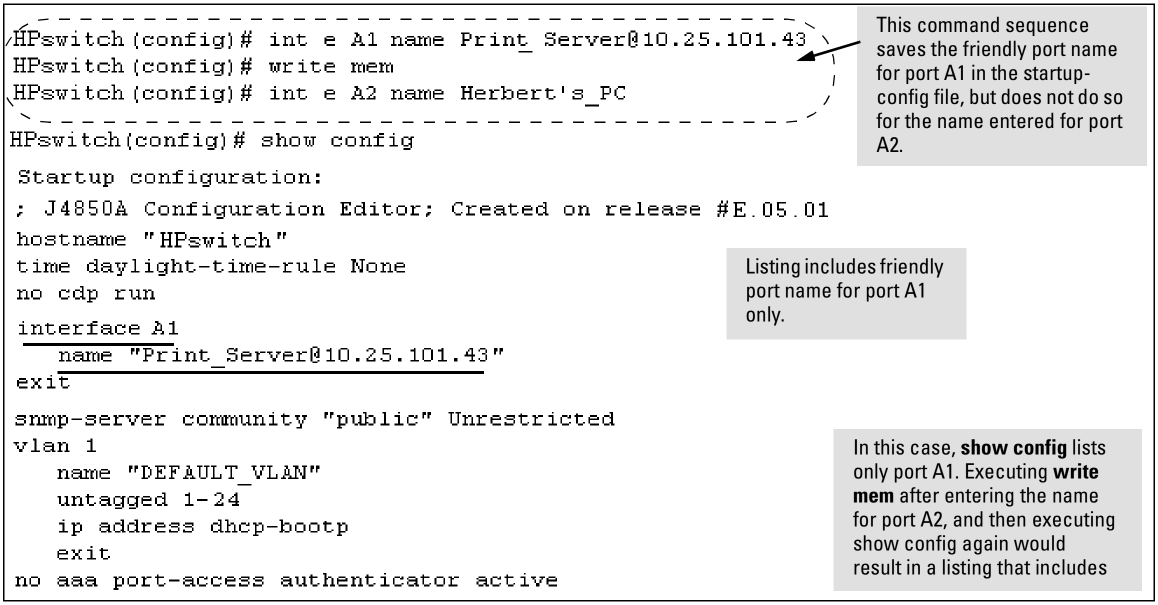

```
HPswitch(config)# int e A1 name Print_Server@10.25.101.43
HPswitch(config)# write mem
HPswitch(config)# int e A2 name Herbert's_PC
HPswitch(config)# show config
 Startup configuration:
 ; J4850A Configuration Editor; Created on release #E.05.01
 hostname "HPswitch"
 time daylight-time-rule None
 no cdp run
 interface A1
    name "Print_Server@10.25.101.43"
 exit
 snmp-server community "public" Unrestricted
 vlan 1
    name "DEFAULT_VLAN"
    untagged 1-24
    ip address dhcp-bootp
    exit
 no aaa port-access authenticator active
```

This command sequence saves the friendly port name for port A1 in the startup-config file, but does not do so for the name entered for port A2.

Listing includes friendly port name for port A1 only.

In this case, **show config** lists only port A1. Executing **write mem** after entering the name for port A2, and then executing show config again would result in a listing that includes

Figure 6-14. Example Listing of the Startup-Config File with a Friendly Port Name Configured (and Saved)

Configuring IP Addressing

Contents

Overview

You can configure IP addressing through all of the switch's interfaces. You can also:

- Easily edit a switch configuration file to allow downloading the file to multiple Series 5300XL switches without overwriting each switch's unique gateway and VLAN 1 IP addressing.
- Assign up to seven secondary IP address to a VLAN (multinetting).

Why Configure IP Addressing? In its factory default configuration, the switch operates as a multiport learning bridge with network connectivity provided by the ports on the switch. However, to enable specific management access and control through your network, you will need IP addressing. Table 7-1 on page 7-12 shows the switch features that depend on IP addressing to operate.

IP Configuration

IP Configuration Features

Feature	Default	Menu	CLI	Web
IP Address and Subnet Mask	DHCP/Bootp	page 7-5	page 7-7	page 7-11
Multiple IP Addresses on a VLAN	n/a	—	page 7-9	—
Default Gateway Address	none	page 7-5	page 7-7	page 7-11
Packet Time-To-Live (TTL)	64 seconds	page 7-5	page 7-7	—
Time Server (Timep)	DHCP	page 7-5	page 7-7	—

IP Address and Subnet Mask. Configuring the switch with an IP address expands your ability to manage the switch and use its features. By default, the switch is configured to automatically receive IP addressing on the default VLAN from a DHCP/Bootp server that has been configured correctly with information to support the switch. (Refer to “DHCP/Bootp Operation” on page 7-13 for information on setting up automatic configuration from a server.) However, if you are not using a DHCP/Bootp server to configure IP addressing, use the menu interface or the CLI to manually configure the initial IP values. After you have network access to a device, you can use the web browser interface to modify the initial IP configuration if needed.

For information on how IP addressing affects switch performance, refer to “How IP Addressing Affects Switch Operation” on page 7-12.

Multinetting: Assigning Multiple IP Addresses to a VLAN. For a given VLAN you can assign one primary IP address and up to seven secondary IP addresses. This allows you to combine two or more subnets on the same VLAN, which enables devices in the combined subnets to communicate normally through the network without needing to reconfigure the IP addressing in any of the combined subnets.

Default Gateway Operation. The default gateway is required when a router is needed for tasks such as reaching off-subnet destinations or forwarding traffic across multiple VLANs. The gateway value is the IP address of the next-hop gateway node for the switch, which is used if the requested destination address is not on a local subnet/VLAN. If the switch does not have a manually-configured default gateway and DHCP/Bootp is configured on the primary VLAN, then the default gateway value provided by the DHCP or Bootp server will be used. If the switch has a manually configured default gateway, then the switch uses this gateway, even if a different gateway is received via DHCP or Bootp on the primary VLAN. (This is also true for TimeP and a non-default Time-To-Live.) See "Notes" on page 7-4 and "The Primary VLAN" on page 11-6.

Packet Time-To-Live (TTL). This parameter specifies the maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet's TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it. In most cases, the default setting (64) is adequate.

Just Want a Quick Start with IP Addressing?

If you just want to give the switch an IP address so that it can communicate on your network, or if you are not using VLANs, HP recommends that you use the Switch Setup screen to quickly configure IP addressing. To do so, do one of the following:

- Enter `setup` at the CLI Manager level prompt.

  ```
  HPswitch# setup
  ```

- Select **8. Run Setup** in the Main Menu of the menu interface.

For more on using the Switch Setup screen, see the *Installation and Getting Started Guide* you received with the switch.

IP Addressing with Multiple VLANs

In the factory-default configuration, the switch has one, permanent default VLAN (named DEFAULT_VLAN) that includes all ports on the switch. Thus, when only the default VLAN exists in the switch, if you assign an IP address and subnet mask to the switch, you are actually assigning the IP addressing to the DEFAULT_VLAN.

Notes

- If multiple VLANs are configured, then each VLAN can have its own IP address. This is because each VLAN operates as a separate broadcast domain and requires a unique IP address and subnet mask. A default gateway (IP) address for the switch is optional, but recommended.
- In the factory-default configuration, the default VLAN (named DEFAULT_VLAN) is the switch's *primary* VLAN. The switch uses the primary VLAN for learning the default gateway address, (packet) Time-To-Live (TTL), and Timep via DHCP or Bootp. (Other VLANs can also use DHCP or BootP to acquire IP addressing. However, the switch's gateway,

TTL, and TimeP values will be acquired through the primary VLAN only.) For more on VLANs, refer to "Port-Based Virtual LANs (Static VLANs)" on page 11-3.

- The IP addressing used in the switch should be compatible with your network. That is, the IP address must be unique and the subnet mask must be appropriate for your IP network.
- If you plan to connect to other networks that use globally administered IP addresses, refer to "Globally Assigned IP Network Addresses" on page 7-21.
- If you change the IP address through either Telnet access or the web browser interface, the connection to the switch will be lost. You can reconnect by either restarting Telnet with the new IP address or entering the new address as the URL in your web browser.

Menu: Configuring IP Address, Gateway, and Time-To-Live (TTL)

Do one of the following:

- To manually enter an IP address, subnet mask, set the **IP Config** parameter to **Manual** and then manually enter the IP address and subnet mask values you want for the switch.
- To use DHCP or Bootp, use the menu interface to ensure that the **IP Config** parameter is set to **DHCP/Bootp**, then refer to "DHCP/Bootp Operation" on page 7-13.

To Configure IP Addressing.

1. From the Main Menu, Select.

   2. Switch Configuration ...

      5. IP Configuration

Notes

If multiple VLANs are configured, a screen showing all VLANs appears instead of the following screen.

The Menu interface displays only the primary IP address for any VLAN. If you use the CLI to configure secondary IP addresses on a VLAN, use the CLI **show ip** command to list them. (Refer to "Viewing the Current IP Configuration" on page 7-7.)

For descriptions of these parameters, see the online Help for this screen.

Before using the DHCP/Bootp option, refer to "DHCP/Bootp Operation" on page 7-13.

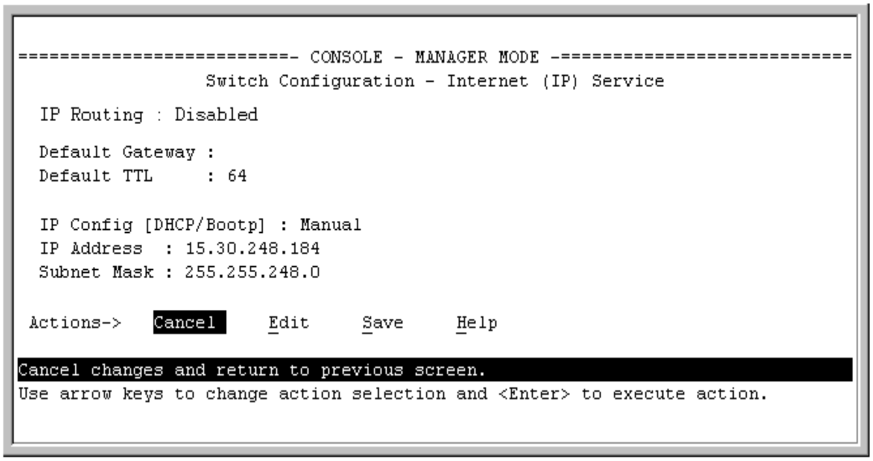

```
==========================- CONSOLE - MANAGER MODE -============================
                 Switch Configuration - Internet (IP) Service

  IP Routing : Disabled

  Default Gateway :
  Default TTL     : 64

  IP Config [DHCP/Bootp] : Manual
  IP Address  : 15.30.248.184
  Subnet Mask : 255.255.248.0

 Actions->  Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 7-1. Example of the IP Service Configuration Screen without Multiple VLANs Configured

2. Press [E] (for **Edit**).
3. If the switch needs to access a router, for example, to reach off-subnet destinations, select the **Default Gateway** field and enter the IP address of the gateway router.
4. If you need to change the packet Time-To-Live (TTL) setting, select **Default TTL** and type in a value between 2 and 255.
5. To configure IP addressing, select **IP Config** and do one of the following:
   - If you want to have the switch retrieve its IP configuration from a DHCP or Bootp server, at the **IP Config** field, keep the value as **DHCP/Bootp** and go to step 8.
   - If you want to manually configure the IP information, use the Space bar to select **Manual** and use the [Tab] key to move to the other IP configuration fields.
6. Select the **IP Address** field and enter the IP address for the switch.
7. Select the **Subnet Mask** field and enter the subnet mask for the IP address.
8. Press [Enter], then [S] (for **Save**).

CLI: Configuring IP Address, Gateway, and Time-To-Live (TTL)

IP Commands Used in This Section	Page
show ip	7-7
show ip secondary	7-10
ip address /< *mask-length* | *mask-bits* >	7-8
ip address /< *mask-length* | *mask-bits* > secondary	7-9
ip default-gateway	7-11
ip ttl	7-11

Viewing the Current IP Configuration. The following command displays the IP addressing for each VLAN configured in the switch. If only the DEFAULT_VLAN exists, then its IP configuration applies to all ports in the switch. Where multiple VLANs are configured, the IP addressing is listed per VLAN. The display includes switch-wide packet time-to-live, and (if configured) the switch's default gateway and Timep configuration.

Syntax: show ip

For example, in the factory-default configuration (no IP addressing assigned), the switch's IP addressing appears as:

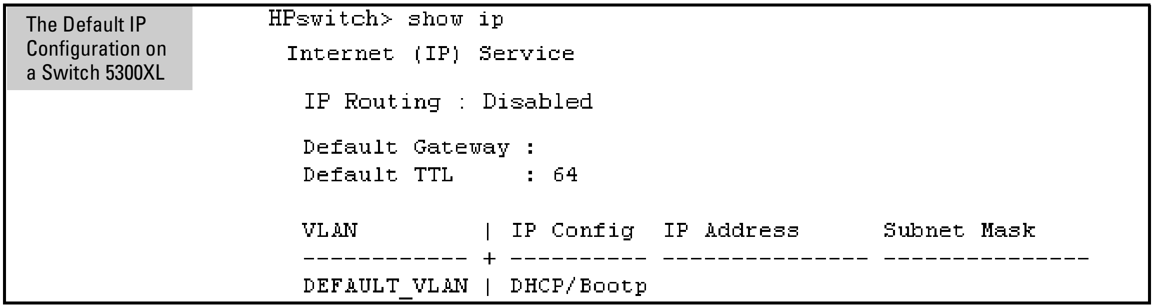

The Default IP Configuration on a Switch 5300XL

```
HPswitch> show ip
 Internet (IP) Service

  IP Routing : Disabled

  Default Gateway :
  Default TTL     : 64

  VLAN         | IP Config  IP Address      Subnet Mask
  ------------ + ---------- --------------- ---------------
  DEFAULT_VLAN | DHCP/Bootp
```

Figure 7-2. Example of the Switch's Default IP Addressing

With multiple VLANs and some other features configured, **show ip** provides additional information:

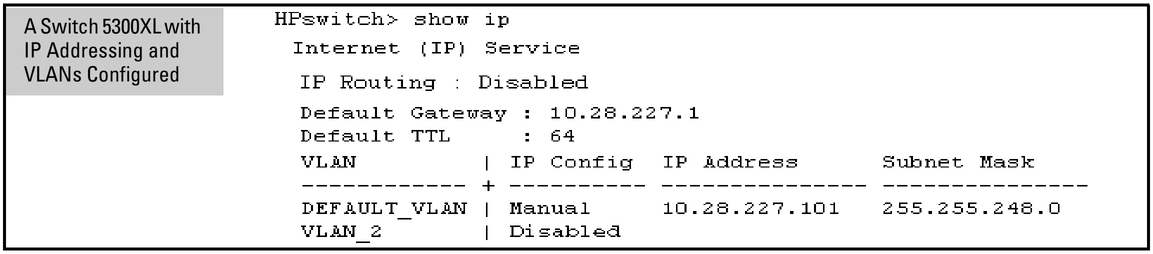

A Switch 5300XL with IP Addressing and VLANs Configured

```
HPswitch> show ip
 Internet (IP) Service
  IP Routing : Disabled
  Default Gateway : 10.28.227.1
  Default TTL     : 64
  VLAN         | IP Config  IP Address      Subnet Mask
  ------------ + ---------- --------------- ---------------
  DEFAULT_VLAN | Manual     10.28.227.101   255.255.248.0
  VLAN_2       | Disabled
```

Figure 7-3. Example of Show IP Listing with Non-Default IP Addressing Configured

Configure an IP Address and Subnet Mask. The following command includes both the primary IP address and the subnet mask. You must either include the ID of the VLAN for which you are configuring IP addressing or go to the context configuration level for that VLAN. (If you are not using VLANs on the switch—that is, if the only VLAN is the default VLAN—then the VLAN ID is always "1".)

Note

The default IP address setting for the DEFAULT_VLAN is **DHCP/Bootp**. On additional VLANs you create, the default IP address setting is **Disabled**.

Syntax: [no] vlan < *vlan-id* > ip address <*ip-address/mask-length*>
or
[no] vlan < *vlan-id* > ip address < *ip-address* > < *mask-bits* >
or
vlan < *vlan-id* > ip address *dhcp-bootp*

This example configures IP addressing on the default VLAN with the subnet mask specified in mask bits.

```
HPswitch(config)# vlan 1 ip address 10.28.227.103 255.255.255.0
```

This example configures the same IP addressing as the preceding example, but specifies the subnet mask by mask length.

```
HPswitch(config)# vlan 1 ip address 10.28.227.103/24
```

This example deletes an IP address configured in VLAN 1.

```
HPswitch (config) no vlan 1 ip address 10.28.227.103/24
```

Configure Multiple IP Addresses on a VLAN (Multinetting). You can configure one primary IP address per VLAN and up to seven secondary IP addresses for the same VLAN. That is, the switch enables you to assign up to eight networks to a VLAN.

- Each IP address on a VLAN must be for a separate subnet.
- The VLAN must have a manual primary IP address configured before any secondary addresses configured on the VLAN will be enabled.
- If the primary IP address on a VLAN is configured for DHCP-Bootp, the switch does not accept secondary IP addresses on that VLAN.
- DHCP operates only to provide primary IP addressing, and is not used for providing secondary IP addressing.
- The switch allows up to 512 secondary subnet address assignments to VLANs.

Syntax: [no] vlan *<vlan-id>* ip address *<ip-address/mask-length> secondary*
[no] vlan *<vlan-id>* ip address *<ip-address> <mask-bits> secondary*

For example, if you wanted to multinet VLAN_20 (VID = 20) with its primary IP address and two secondary IP addresses shown below, you would perform steps similar to the following. (For this example, assume that the primary IP addressing is already configured.)

Status	VID	IP Address	Subnet Mask
Primary	20	10.25.33.101	255.255.240.0
Secondary	20	10.26.33.101	255.255.240.0
Secondary	20	10.27.33.101	255.255.240.0

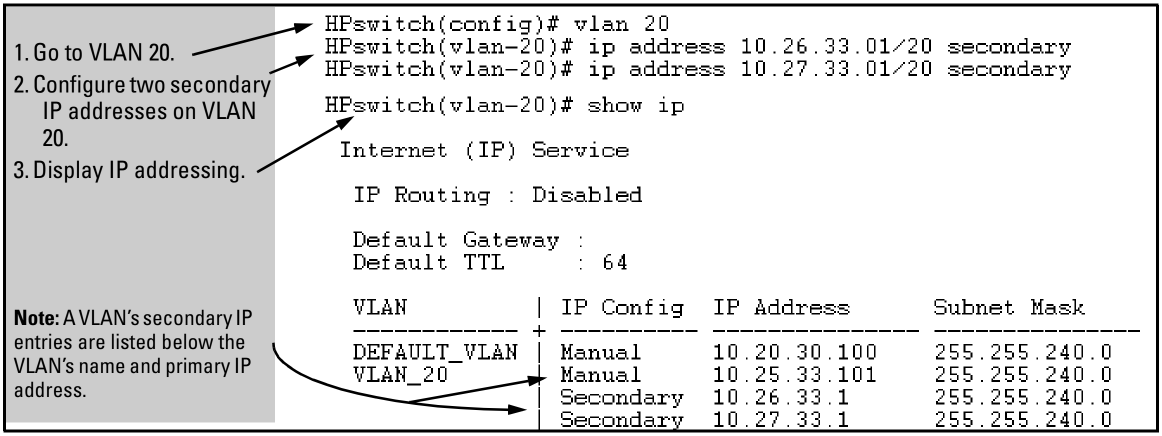

1. Go to VLAN 20.
2. Configure two secondary IP addresses on VLAN 20.
3. Display IP addressing.

Note: A VLAN's secondary IP entries are listed below the VLAN's name and primary IP address.

```
HPswitch(config)# vlan 20
HPswitch(vlan-20)# ip address 10.26.33.01/20 secondary
HPswitch(vlan-20)# ip address 10.27.33.01/20 secondary
HPswitch(vlan-20)# show ip

 Internet (IP) Service

  IP Routing : Disabled

  Default Gateway :
  Default TTL     : 64

  VLAN         | IP Config  IP Address      Subnet Mask
  ------------ + ---------- --------------- ---------------
  DEFAULT_VLAN | Manual     10.20.30.100    255.255.240.0
  VLAN_20      | Manual     10.25.33.101    255.255.240.0
               | Secondary  10.26.33.1      255.255.240.0
               | Secondary  10.27.33.1      255.255.240.0
```

Figure 7-4. Example of Configuring and Displaying a Multinetted VLAN

If you then wanted to multinet the default VLAN, you would do the following:

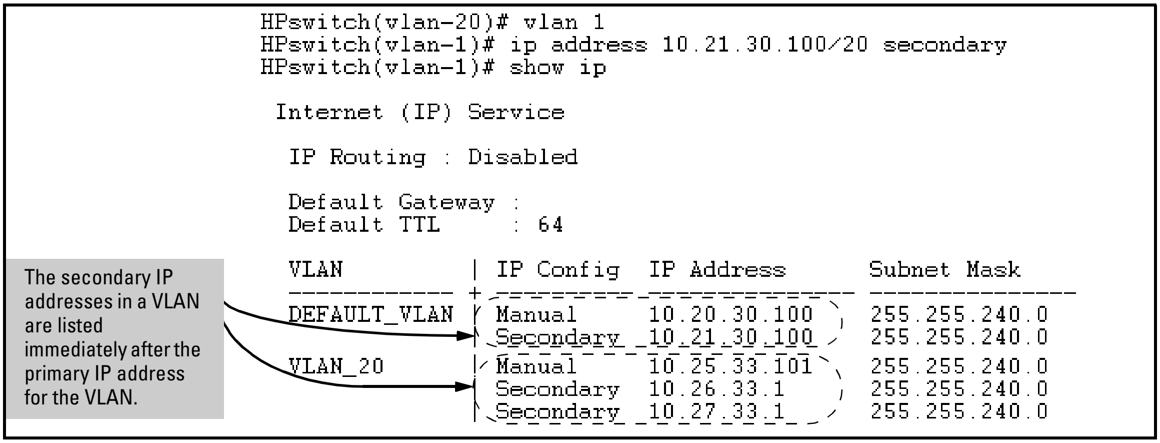

```
HPswitch(vlan-20)# vlan 1
HPswitch(vlan-1)# ip address 10.21.30.100/20 secondary
HPswitch(vlan-1)# show ip

 Internet (IP) Service

  IP Routing : Disabled

  Default Gateway :
  Default TTL     : 64

  VLAN         | IP Config  IP Address      Subnet Mask
  ------------ + ---------- --------------- ---------------
  DEFAULT_VLAN | Manual     10.20.30.100    255.255.240.0
               | Secondary  10.21.30.100    255.255.240.0
  VLAN_20      | Manual     10.25.33.101    255.255.240.0
               | Secondary  10.26.33.1      255.255.240.0
               | Secondary  10.27.33.1      255.255.240.0
```

The secondary IP addresses in a VLAN are listed immediately after the primary IP address for the VLAN.

Figure 7-5. Example of Multinetting on the Default VLAN

For an alternate display of the above information, use the **show ip secondary** command:

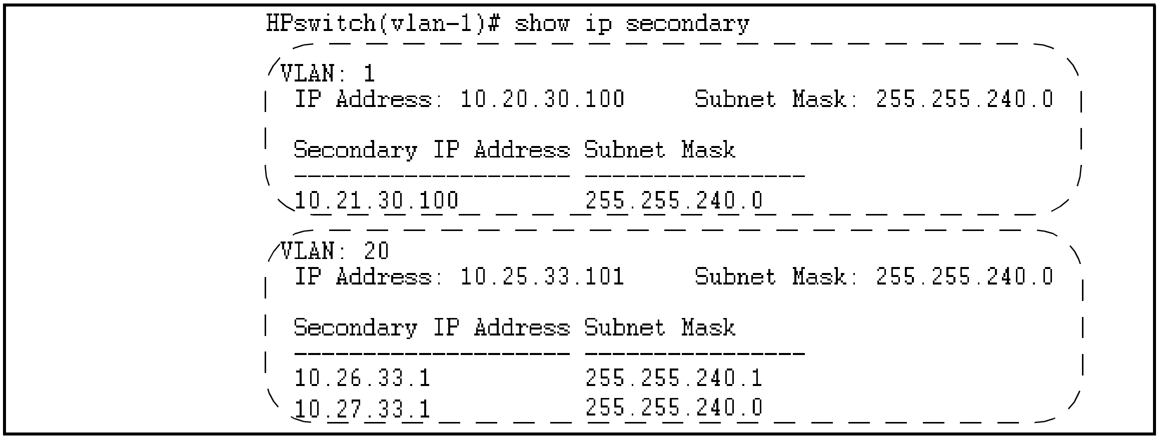

```
HPswitch(vlan-1)# show ip secondary

VLAN: 1
  IP Address: 10.20.30.100     Subnet Mask: 255.255.240.0

  Secondary IP Address Subnet Mask
  -------------------- ---------------
  10.21.30.100         255.255.240.0

VLAN: 20
  IP Address: 10.25.33.101     Subnet Mask: 255.255.240.0

  Secondary IP Address Subnet Mask
  -------------------- ---------------
  10.26.33.1           255.255.240.1
  10.27.33.1           255.255.240.0
```

Figure 7-6. Show IP Secondary Lists Secondary IP Groups

Configure the Optional Default Gateway. Using the Global configuration level, you can assign one default gateway to the switch.

Syntax: ip default-gateway <*ip-address*>

For example:

```
HPswitch(config)# ip default-gateway 10.28.227.115
```

Note

The switch uses the IP default gateway only while operating as a Layer 2 device. While routing is enabled on the switch, the IP default gateway is not used. Thus, to avoid loss of Telnet access to off-subnet management stations, you should use the **ip route** command to configure a static (default) route before enabling routing. Refer to chapter 16, "IP Routing Features", for more information.

Configure Time-To-Live (TTL). The maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet's TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it.

Syntax: ip ttl <*number-of-hops*>

```
HPswitch(config)# ip ttl 60
```

In the CLI, you can execute this command only from the global configuration level. The TTL default is 64, and the range is 2 - 255.

Web: Configuring IP Addressing

You can use the web browser interface to access IP addressing only if the switch already has an IP address that is reachable through your network.

1. Click on the Configuration tab.
2. Click on [IP Configuration].
3. If you need further information on using the web browser interface, click on [?] to access the web-based help available for the Series 5300XL switches.

How IP Addressing Affects Switch Operation

Without an IP address and subnet mask compatible with your network, the switch can be managed only through a direct terminal device connection to the Console RS-232 port. You can use direct-connect console access to take advantage of features that do not depend on IP addressing. However, to realize the full performance capabilities HP proactive networking offers through the switch, configure the switch with an IP address and subnet mask compatible with your network. The following table lists the general features available with and without a network-compatible IP address configured.

Table 7-1. Features Available With and Without IP Addressing on the Switch

Features Available Without an IP Address	Additional Features Available with an IP Address and Subnet Mask
• Direct-connect access to the CLI and the menu interface. • DHCP or Bootp support for automatic IP address configuration, and DHCP support for automatic Timep server IP address configuration • Spanning Tree Protocol • Port settings and port trunking • Console-based status and counters information for monitoring switch operation and diagnosing problems through the CLI or menu interface. • VLANs and GVRP • Serial downloads of operating system (OS) updates and configuration files (Xmodem) • Link test • Port monitoring • Password authentication • Quality of Service (QoS) • Authorized IP manager security	• HP web browser interface access, with configuration, security, and diagnostic tools, plus the Alert Log for discovering problems detected in the switch along with suggested solutions • SNMP network management access such as HP TopTools network configuration, monitoring, problem-finding and reporting, analysis, and recommendations for changes to increase control and uptime • TACACS+, RADIUS, SSH, and 802.1x authentication • Multinetting on VLANs • CDP support • Telnet access to the CLI or the menu interface • IGMP • Timep server configuration • TFTP download of configurations and OS updates • Ping test

DHCP/Bootp Operation

Overview. DHCP/Bootp is used to provide configuration data from a DHCP or Bootp server to the switch. This data can be the IP address, subnet mask, default gateway, Timep Server address, and TFTP server address. If a TFTP server address is provided, this allows the switch to TFTP a previously saved configuration file from the TFTP server to the switch. With either DHCP or Bootp, the servers must be configured prior to the switch being connected to the network.

Note

The Series 5300XL switches are compatible with both DHCP and Bootp servers.

The DHCP/Bootp Process. Whenever the **IP Config** parameter in the switch or in an individual VLAN in the switch is configured to **DHCP/Bootp** (the default), or when the switch is rebooted with this configuration:

1. DHCP/Bootp requests are automatically broadcast on the local network. (The switch sends one type of request to which either a DHCP or Bootp server can respond.)
2. When a DHCP or Bootp server receives the request, it replies with a previously configured IP address and subnet mask for the switch. The switch also receives an IP Gateway address if the server has been configured to provide one. In the case of Bootp, the server must first be configured with an entry that has the switch's MAC address. (To determine the switch's MAC address, see appendix D, "MAC Address Management".) The switch properly handles replies from either type of server. If multiple replies are returned, the switch tries to use the first reply.)

Note

If you manually configure a gateway on the switch, it will ignore any gateway address received via DHCP or Bootp.

If the switch is initially configured for DHCP/Bootp operation (the default), or if it reboots with this configuration, it begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch may not immediately receive the desired configuration. After verifying that the server has become accessible to the switch, reboot the switch to re-start the process immediately.

DHCP Operation. A significant difference between a DHCP configuration and a Bootp configuration is that an IP address assignment from a DHCP server is automatic. Depending on how the DHCP server is configured, the switch may receive an ip address that is temporarily leased. Periodically the switch may be required to renew its lease of the IP configuration. Thus, the IP addressing provided by the server may be different each time the switch reboots or renews its configuration from the server. However, you can fix the address assignment for the switch by doing either of the following:

- Configure the server to issue an “infinite” lease.
- Using the switch’s MAC address as an identifier, configure the server with a “Reservation” so that it will always assign the same IP address to the switch. (For MAC address information, refer to appendix D, “MAC Address Management”.)

For more information on either of these procedures, refer to the documentation provided with the DHCP server.

Bootp Operation. When a Bootp server receives a request it searches its Bootp database for a record entry that matches the MAC address in the Bootp request from the switch. If a match is found, the configuration data in the associated database record is returned to the switch. For many Unix systems, the Bootp database is contained in the **/etc/bootptab** file. In contrast to DHCP operation, Bootp configurations are always the same for a specific receiving device. That is, the Bootp server replies to a request with a configuration previously stored in the server and designated for the requesting device.

Bootp Database Record Entries. A minimal entry in the Bootp table file **/etc/bootptab** to update an IP address and subnet mask to the switch or a VLAN configured in the switch would be similar to this entry:

```
5300switch:\
   ht=ether:\
   ha=0030c1123456:\
   ip=10.66.77.88:\
   sm=255.255.248.0:\
   gw=10.66.77.1:\
   hn:\
   vm=rfc1048
```

An entry in the Bootp table file `/etc/bootptab` to tell the switch or VLAN where to obtain a configuration file download would be similar to this entry:

```
5300switch:\
   ht=ether:\
   ha=0030c1123456:\
   ip=10.66.77.88:\
   sm=255.255.248.0:\
   gw=10.66.77.1:\
   lg=10.22.33.44:\
   T144=”switch.cfg”:\
   vm=rfc1048
```

where:

5300switch — is a user-defined symbolic name to help you find the correct section of the bootptab file. If you have multiple switches that will be using Bootp to get their IP configuration, you should use a unique symbolic name for each switch.

ht — is the “hardware type”. For the Series 5300XL switches, set this to **ether** (for Ethernet). *This tag must precede the* ha *tag.*

ha — is the “hardware address”. Use the switch's (or VLAN's) 12-digit MAC address.

ip — is the IP address to be assigned to the switch (or VLAN).

sm — is the subnet mask of the subnet in which the switch (or VLAN) is installed.

gw — is the IP address of the default gateway.

lg — TFTP server address (source of final configuration file)

T144 — is the vendor-specific “tag” identifying the configuration file to download.

vm — is a required entry that specifies the Bootp report format. For theSeries 5300XL switches, set this parameter to **rfc1048**.

Note

The above Bootp table entry is a sample that will work for the Series 5300XL switches when the appropriate addresses and file names are used.

Network Preparations for Configuring DHCP/Bootp

In its default configuration, the switch is configured for DHCP/Bootp operation. However, the DHCP/Bootp feature will not acquire IP addressing for the switch unless the following tasks have already been completed:

- For Bootp operation:
  - A Bootp database record has already been entered into an appropriate Bootp server.
  - The necessary network connections are in place
  - The Bootp server is accessible from the switch
- For DHCP operation:
  - A DHCP scope has been configured on the appropriate DHCP server.
  - The necessary network connections are in place
  - A DHCP server is accessible from the switch

Note

Designating a primary VLAN other than the default VLAN affects the switch's use of information received via DHCP/Bootp. For more on this topic, see "The Primary VLAN" on page 11-6.

After you reconfigure or reboot the switch with DHCP/Bootp enabled in a network providing DHCP/Bootp service, the switch does the following:

- Receives an IP address and subnet mask and, if configured in the server, a gateway IP address and the address of a Timep server.
- If the DHCP/Bootp reply provides information for downloading a configuration file, the switch uses TFTP to download the file from the designated source, then reboots itself. (This assumes that the switch or VLAN has connectivity to the TFTP file server specified in the reply, that the configuration file is correctly named, and that the configuration file exists in the TFTP directory.)

IP Preserve: Retaining VLAN-1 IP Addressing Across Configuration File Downloads

IP Preserve enables you to copy a configuration file to multiple Series 5300XL switches while retaining the individual IP address and subnet mask on VLAN 1 in each switch, and the Gateway IP address assigned to the switch. This enables you to distribute the same configuration file to multiple switches without overwriting their individual IP addresses.

Operating Rules for IP Preserve

When **ip preserve** is entered as the last line in a configuration file stored on a TFTP server:

- If the switch's current IP address for VLAN 1 was not configured by DHCP/ Bootp, IP Preserve retains the switch's current IP address, subnet mask, and IP gateway address when the switch downloads the file and reboots. The switch adopts all other configuration parameters in the configuration file into the startup-config file.
- If the switch's current IP addressing for VLAN 1 is from a DHCP server, IP Preserve is suspended. In this case, whatever IP addressing the configuration file specifies is implemented when the switch downloads the file and reboots. If the file includes DHCP/Bootp as the IP addressing source for VLAN 1, the switch will configure itself accordingly and use DHCP/ Bootp. If instead, the file includes a dedicated IP address and subnet mask for VLAN 1 and a specific gateway IP address, then the switch will implement these settings in the startup-config file.
- The **ip preserve** statement does not appear in **show config** listings. To verify IP Preserve in a configuration file, open the file in a text editor and view the last line. For an example of implementing IP Preserve in a configuration file, see figure 7-7, below.

To set up IP Preserve, enter the **ip preserve** statement at the end of a configuration file. (Note that you do not execute IP Preserve by entering a command from the CLI).

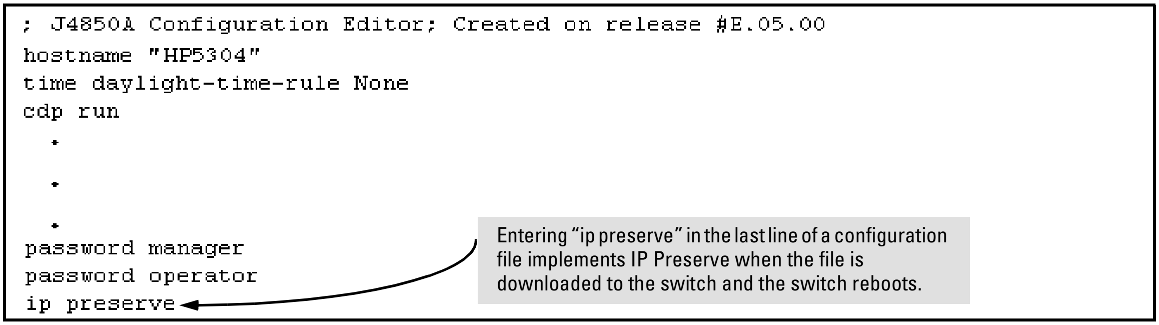

```
; J4850A Configuration Editor; Created on release #E.05.00
hostname "HP5304"
time daylight-time-rule None
cdp run
  .
  .
  .
password manager
password operator
ip preserve
```

Entering "ip preserve" in the last line of a configuration file implements IP Preserve when the file is downloaded to the switch and the switch reboots.

Figure 7-7. Example of Implementing IP Preserve in a Series 5300XL switch Configuration File

For example, consider Figure 7-8:

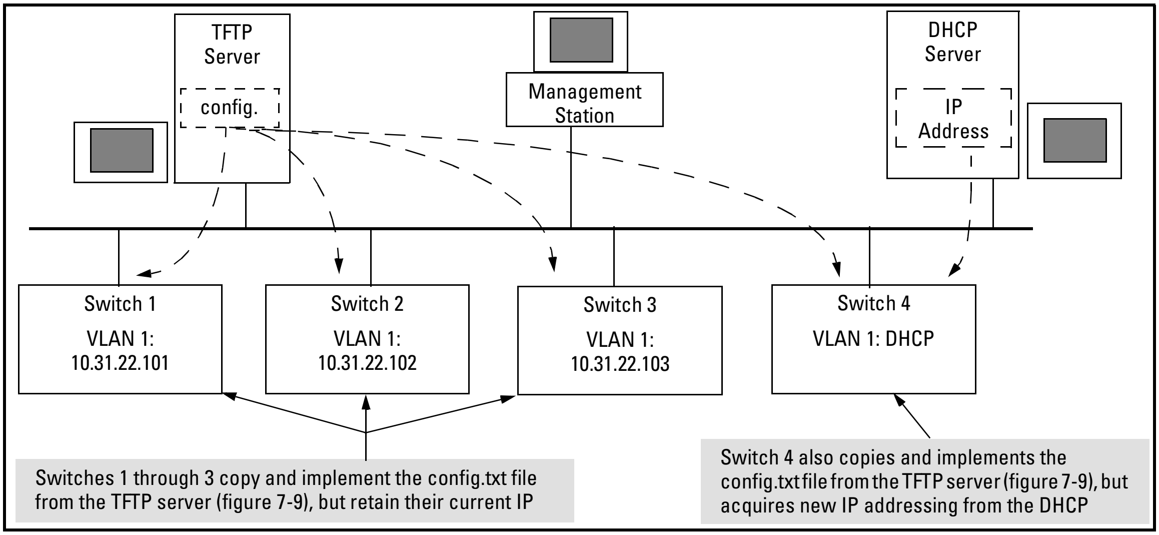

Switches 1 through 3 copy and implement the config.txt file from the TFTP server (figure 7-9), but retain their current IP

Switch 4 also copies and implements the config.txt file from the TFTP server (figure 7-9), but acquires new IP addressing from the DHCP

Figure 7-8. Example of IP Preserve Operation with Multiple Series 5300XL Switches

If you apply the following configuration file to figure 7-8, switches 1 - 3 will retain their manually assigned IP addressing and switch 4 will be configured to acquire its IP addressing from a DHCP server.

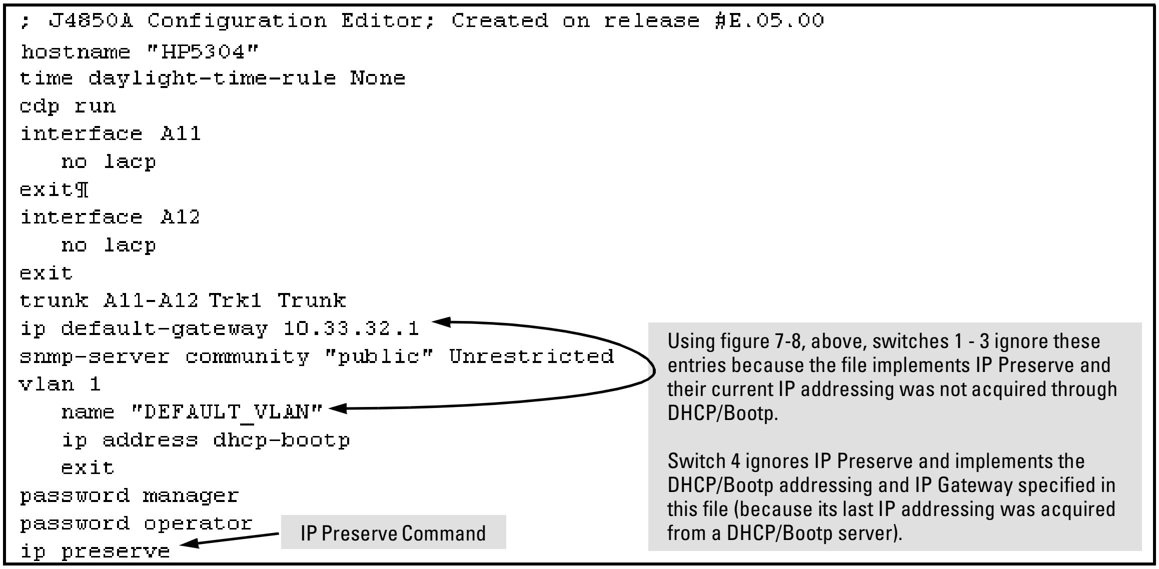

```
; J4850A Configuration Editor; Created on release #E.05.00
hostname "HP5304"
time daylight-time-rule None
cdp run
interface A11
   no lacp
exit¶
interface A12
   no lacp
exit
trunk A11-A12 Trk1 Trunk
ip default-gateway 10.33.32.1
snmp-server community "public" Unrestricted
vlan 1
   name "DEFAULT_VLAN"
   ip address dhcp-bootp
   exit
password manager
password operator
ip preserve
```

Using figure 7-8, above, switches 1 - 3 ignore these entries because the file implements IP Preserve and their current IP addressing was not acquired through DHCP/Bootp.

Switch 4 ignores IP Preserve and implements the DHCP/Bootp addressing and IP Gateway specified in this file (because its last IP addressing was acquired from a DHCP/Bootp server).

IP Preserve Command

Figure 7-9. Configuration File in TFTP Server, with DHCP/Bootp Specified as the IP Addressing Source

If you apply this configuration file to figure 7-8, switches 1 - 3 will still retain their manually assigned IP addressing. However, switch 4 will be configured with the IP addressing included in the file.

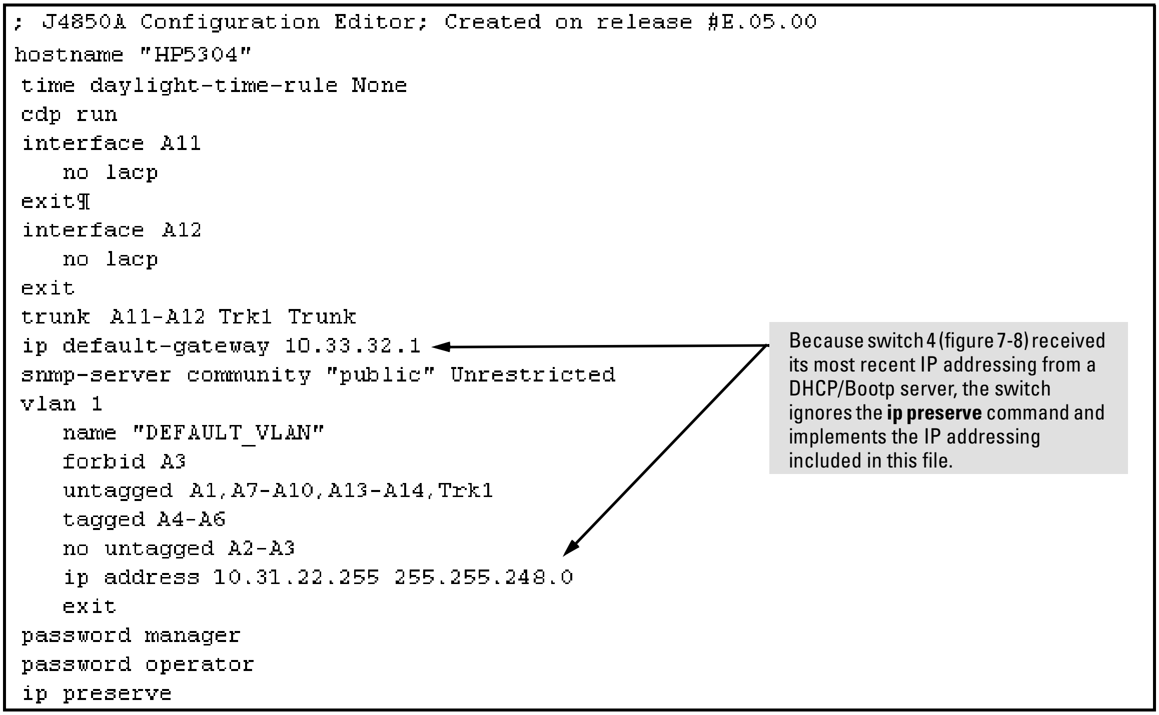

```
; J4850A Configuration Editor; Created on release #E.05.00
hostname "HP5304"
 time daylight-time-rule None
 cdp run
 interface A11
    no lacp
 exit¶
 interface A12
    no lacp
 exit
 trunk A11-A12 Trk1 Trunk
 ip default-gateway 10.33.32.1
 snmp-server community "public" Unrestricted
 vlan 1
    name "DEFAULT_VLAN"
    forbid A3
    untagged A1,A7-A10,A13-A14,Trk1
    tagged A4-A6
    no untagged A2-A3
    ip address 10.31.22.255 255.255.248.0
    exit
 password manager
 password operator
 ip preserve
```

Because switch 4 (figure 7-8) received its most recent IP addressing from a DHCP/Bootp server, the switch ignores the **ip preserve** command and implements the IP addressing included in this file.

Figure 7-10. Configuration File in TFTP Server, with Dedicated IP Addressing Instead of DHCP/Bootp

To summarize the IP Preserve effect on IP addressing:

- If the switch received its most recent VLAN 1 IP addressing from a DHCP/Bootp server, it ignores the IP Preserve command when it downloads the configuration file, and implements whatever IP addressing instructions are in the configuration file.
- If the switch did not receive its most recent VLAN 1 IP addressing from a DHCP/Bootp server, it retains its current IP addressing when it downloads the configuration file.
- The content of the downloaded configuration file determines the IP addresses and subnet masks for other VLANs.

Globally Assigned IP Network Addresses

If you intend to connect your network to other networks that use globally administered IP addresses, Hewlett-Packard strongly recommends that you use IP addresses that have a network address assigned to you. There is a formal process for assigning unique IP addresses to networks worldwide. For more information:

Please contact your internet service provider (ISP).

If you need more information than your ISP can provide, contact one of the following organizations:

Country	**Phone Number/E-Mail/URL**	**Organization Name/Address**
United States/ Countries not in Europe or Asia/Pacific	1-310-823-9358 icann@icann.org http://www.icann.org	The Internet Corporation for Assigned Names and Numbers (ICANN) 4676 Admiralty Way, Suite 330 Marina Del Rey, CA 90292 USA
Europe	+31 20 535 4444 ncc@ripe.net http://www.ripe.net	RIPE NCC Singel 258 1016 AB Amsterdam The Netherlands
Asia/Pacific	+61-7-3367-0490 info@apnic.net http://www.apnic.net	Attention: IN-ADDR.ARPA Registration Asia Pacific Network Information Center Level 1, 33 Park Road PO Box 2131 Milton, QLD 4064 Australia

For more information, refer to the latest edition of *Internetworking with TCP/IP: Principles, Protocols and Architecture* by Douglas E. Comer (Prentice-Hall, Inc., publisher).

8

Time Protocols

Contents

Overview

This chapter describes:

- SNTP Time Protocol Operation
- Timep Time Protocol Operation

Using time synchronization ensures a uniform time among interoperating devices. This helps you to manage and troubleshoot switch operation by attaching meaningful time data to event and error messages.

The switch offers TimeP and SNTP (Simple Network Time Protocol) and a **timesync** command for changing the time protocol selection (or turning off time protocol operation).

Notes

- Although you can create and save configurations for both time protocols without conflicts, the switch allows only one active time protocol at any time.
- In the factory-default configuration, the time synchronization option is set to TimeP, with the TimeP mode itself set to **Disabled**.

TimeP Time Synchronization

You can either manually assign the switch to use a TimeP server or use DHCP to assign the TimeP server. In either case, the switch can get its time synchronization updates from only one, designated Timep server. This option enhances security by specifying which time server to use.

SNTP Time Synchronization

SNTP provides two operating modes:

- **Broadcast Mode:** The switch acquires time updates by accepting the time value from the first SNTP time broadcast detected. (In this case, the SNTP server must be configured to broadcast time updates to the network broadcast address. Refer to the documentation provided with your SNTP server application.) Once the switch detects a particular server, it ignores time broadcasts from other SNTP servers unless the configurable **Poll Interval** expires three consecutive times without an update received from the first-detected server.

Note

To use Broadcast mode, the switch and the SNTP server must be in the same subnet.

- **Unicast Mode:** The switch requests a time update from the configured SNTP server. (You can configure one server using the menu interface, or up to three servers using the CLI **sntp server** command.) This option provides increased security over the Broadcast mode by specifying which time server to use instead of using the first one detected through a broadcast.

Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation

General Steps for Running a Time Protocol on the Switch:

1. Select the time synchronization protocol: **SNTP** or **TimeP** (the default).
2. Enable the protocol. The choices are:
   - SNTP: **Broadcast** or **Unicast**
   - TimeP: **DHCP** or **Manual**
3. Configure the remaining parameters for the time protocol you selected.

   The switch retains the parameter settings for both time protocols even if you change from one protocol to the other. Thus, if you select a time protocol, the switch uses the parameters you last configured for the selected protocol.

Note that simply selecting a time synchronization protocol does not enable that protocol on the switch unless you also enable the protocol itself (step 2, above). For example, in the factory-default configuration, TimeP is the selected time synchronization method. However, because TimeP is disabled in the factory-default configuration, no time synchronization protocol is running.

Disabling Time Synchronization

You can use either of the following methods to disable time synchronization without changing the Timep or SNTP configuration:

- In the System Information screen of the Menu interface, set the **Time Synch Method** parameter to **None**, then press [Enter], then [S] (for **Save**).
- In the Global config level of the CLI, execute **no timesync**.

SNTP: Viewing, Selecting, and Configuring

SNTP Feature	Default	Menu	CLI	Web
view the SNTP time synchronization configuration	n/a	page 8-5	page 8-8	—
select SNTP as the time synchronization method	timep	page 8-6	page 8-9 ff.	—
disable time synchronization	timep	page 8-6	page 8-12	—
enable the SNTP mode (Broadcast, Unicast, or Disabled)	disabled			—
broadcast	n/a	page 8-6	page 8-9	—
unicast	n/a	page 8-6	page 8-10	—
none/disabled	n/a	page 8-6	page 8-13	—
configure an SNTP server address (for Unicast mode only)	none	page 8-6	page 8-10 ff.	—
change the SNTP server version (for Unicast mode only)	3	page 8-7	page 8-12	—
change the SNTP poll interval	720 seconds	page 8-7	page 8-12	—

Table 8-1.SNTP Parameters

SNTP Parameter	Operation
Time Sync Method	Used to select either SNTP, TIMEP, or None as the time synchronization method.
SNTP Mode	
Disabled	The Default. SNTP does not operate, even if specified by the Menu interface **Time Sync Method** parameter or the CLI **timesync** command.
Unicast	Directs the switch to poll a specific server for SNTP time synchronization. Requires at least one server address.
Broadcast	Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval expires three times without the switch detecting a time update from the original server, it the switch accepts a broadcast time update from the next server it detects.
Poll Interval **(seconds)**	In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update. In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update.
Server Address	Used only when the **SNTP Mode** is set to **Unicast**. Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates. You can configure up to three servers; one using the menu or CLI, and two more using the CLI. See “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 22.
Server Version	Default: 3; range: 1 - 7. Specifies the SNTP software version to use, and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3.

Menu: Viewing and Configuring SNTP

To View, Enable, and Modify SNTP Time Protocol:

1. From the Main Menu, select:

   2. Switch Configuration...

      1. System Information

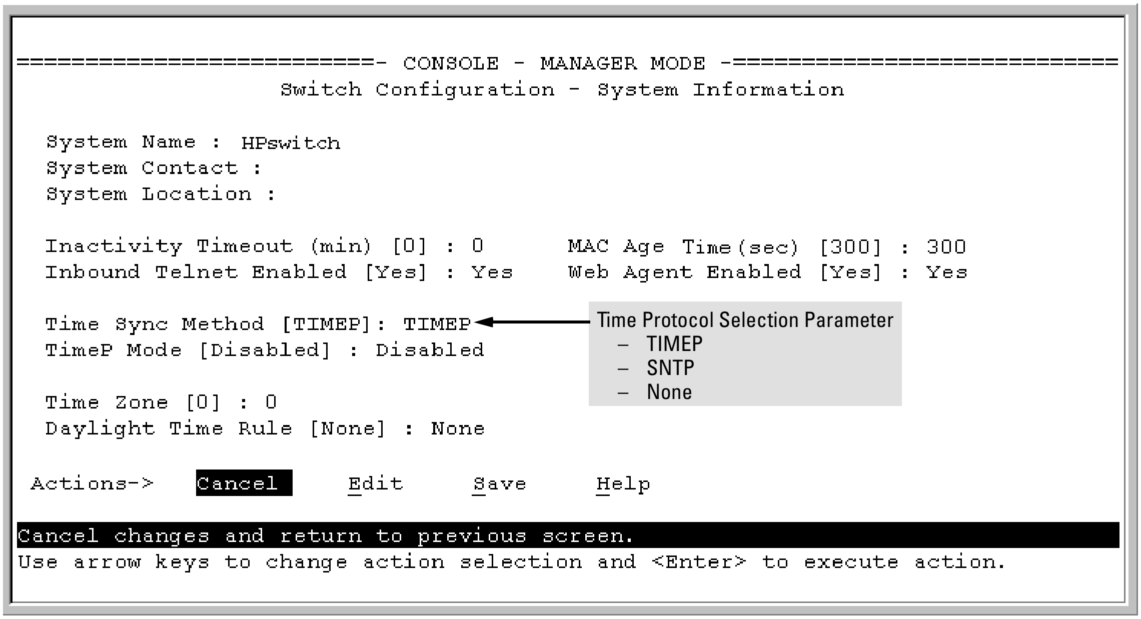

```
=========================- CONSOLE - MANAGER MODE -============================
                   Switch Configuration - System Information

  System Name : HPswitch
  System Contact :
  System Location :

  Inactivity Timeout (min) [0] : 0      MAC Age Time(sec) [300] : 300
  Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes

  Time Sync Method [TIMEP]: TIMEP
  TimeP Mode [Disabled] : Disabled

  Time Zone [0] : 0
  Daylight Time Rule [None] : None

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Time Protocol Selection Parameter
- TIMEP
- SNTP
- None

Figure 8-1. The System Information Screen (Default Values)

2. Press [E] (for **Edit**). The cursor moves to the **System Name** field.
3. Use [↓] to move the cursor to the **Time Sync Method** field.
4. Use the Space bar to select **SNTP**, then press [↓] once to display and move to the **SNTP Mode** field.
5. Do one of the following:
   - Use the Space bar to select the **Broadcast** mode, then press [↓] to move the cursor to the **Poll Interval** field, and go to step 6. (For Broadcast mode details, see "SNTP Operating Modes" on page 8-2.)

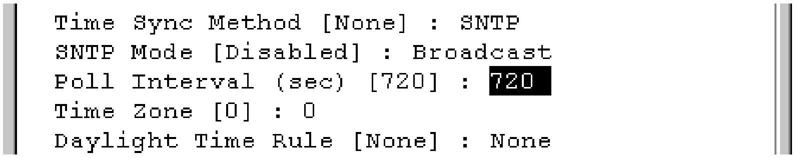

```
Time Sync Method [None] : SNTP
SNTP Mode [Disabled] : Broadcast
Poll Interval (sec) [720] : 720
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

   - Use the Space bar to select the **Unicast** mode, then do the following:
     i. Press [→] to move the cursor to the **Server Address** field.

ii. Enter the IP address of the SNTP server you want the switch to use for time synchronization.

Note: This step replaces any previously configured server IP address. If you will be using backup SNTP servers (requires use of the CLI), then see "SNTP Unicast Time Polling with Multiple SNTP Servers" on page 8-22.

iii. Press [↓] to move the cursor to the **Server Version** field. Enter the value that matches the SNTP server version running on the device you specified in the preceding step (step ii). If you are unsure which version to use, HP recommends leaving this value at the default setting of **3** and testing SNTP operation to determine whether any change is necessary.

Note: Using the menu to enter the IP address for an SNTP server when the switch already has one or more SNTP servers configured causes the switch to delete the primary SNTP server from the server list and to select a new primary SNTP server from the IP address(es) in the updated list. For more on this topic, see "SNTP Unicast Time Polling with Multiple SNTP Servers" on page 8-22.

iv. Press [→] to move the cursor to the **Poll Interval** field, then go to step 6.

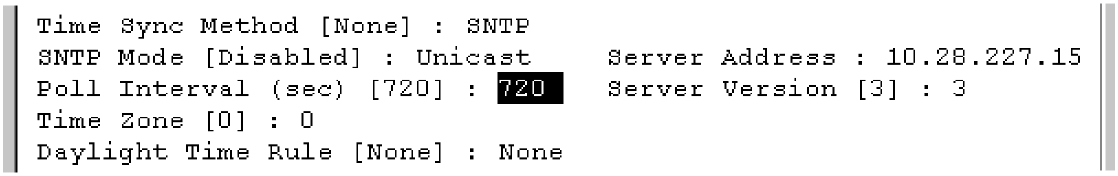

```
Time Sync Method [None] : SNTP
SNTP Mode [Disabled] : Unicast     Server Address : 10.28.227.15
Poll Interval (sec) [720] : 720    Server Version [3] : 3
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

6. In the **Poll Interval** field, enter the time in seconds that you want for a Poll Interval. (For Poll Interval operation, see table 8-1, "SNTP Parameters", on page 8-5.)

7. Press [Enter] to return to the Actions line, then [S] (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

CLI: Viewing and Configuring SNTP

CLI Commands Described in this Section

SNTP Command	Page
show sntp	8-8
[no] timesync	8-9 and ff., 8-12
sntp broadcast	8-9
sntp unicast	8-10
sntp server	8-10 and ff.
Protocol Version	8-12
poll-interval	8-12
no sntp	8-13

This section describes how to use the CLI to view, enable, and configure SNTP parameters.

Viewing the Current SNTP Configuration

This command lists both the time synchronization method (**TimeP**, **SNTP**, or **None**) and the SNTP configuration, even if SNTP is not the selected time protocol.

Syntax: show sntp

For example, if you configured the switch with SNTP as the time synchronization method, then enabled SNTP in broadcast mode with the default poll interval, **show sntp** lists the following:

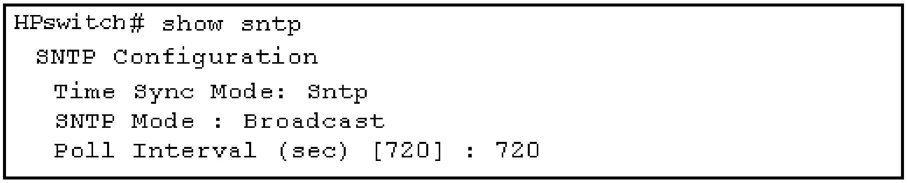

```
HPswitch# show sntp
 SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 720
```

Figure 8-2. Example of SNTP Configuration When SNTP Is the Selected Time Synchronization Method

In the factory-default configuration (where TimeP is the selected time synchronization method), **show sntp** still lists the SNTP configuration even though it is not currently in use. For example:

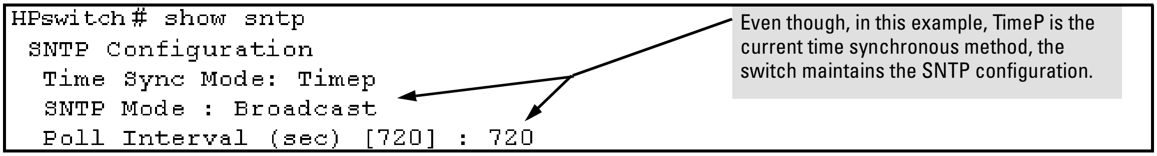

```
HPswitch# show sntp
 SNTP Configuration
  Time Sync Mode: Timep
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 720
```

Even though, in this example, TimeP is the current time synchronous method, the switch maintains the SNTP configuration.

Figure 8-3. Example of SNTP Configuration When SNTP Is Not the Selected Time Synchronization Method

Configuring (Enabling or Disabling) the SNTP Mode

Enabling the SNTP mode means to configure it for either broadcast or unicast mode. Remember that to run SNTP as the switch's time synchronization protocol, you must also select SNTP as the time synchronization method by using the CLI **timesync** command (or the Menu interface **Time Sync Method** parameter).

Syntax:	timesync sntp	Selects SNTP as the time protocol.
	sntp < broadcast | unicast >	Enables the SNTP mode (below and page 8-10).
	sntp server < *ip-addr*>	Required only for unicast mode (page 8-10).
	sntp poll-interval < 30 . . 720>	Enabling the SNTP mode also enables the SNTP poll interval (default: 720 seconds; page 8-12).

Enabling SNTP in Broadcast Mode. Because the switch provides an SNTP polling interval (default: 720 seconds), you need only these two commands for minimal SNTP broadcast configuration:

Syntax:	timesync sntp	Selects SNTP as the time synchronization method.
	sntp broadcast	Configures **Broadcast** as the SNTP mode.

For example, suppose:

- Time synchronization is in the factory-default configuration (TimeP is the currently selected time synchronization method).
- You want to:
  1. View the current time synchronization.
  2. Select SNTP as the time synchronization mode.
  3. Enable SNTP for Broadcast mode.
  4. View the SNTP configuration again to verify the configuration.

The commands and output would appear as follows:

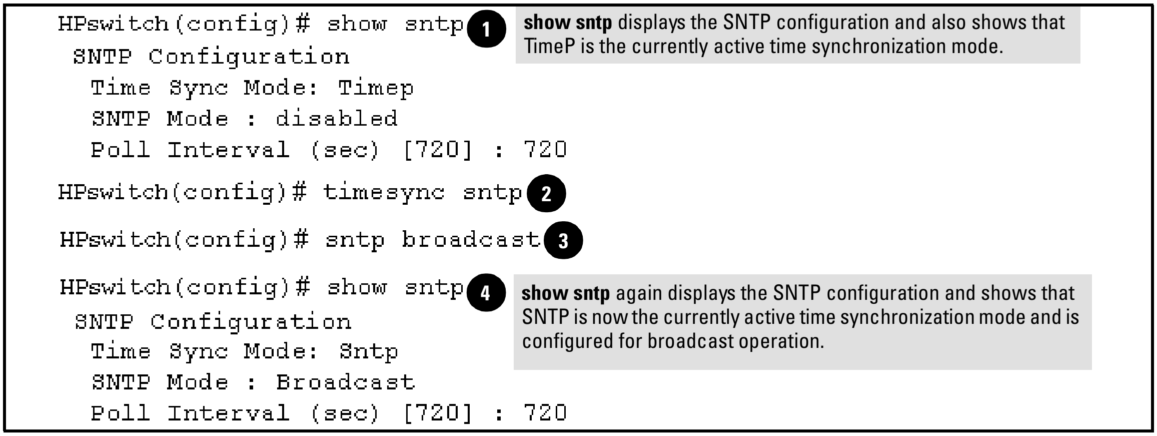

```
HPswitch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Timep
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720
HPswitch(config)# timesync sntp
HPswitch(config)# sntp broadcast
HPswitch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 720
```

1 **show sntp** displays the SNTP configuration and also shows that TimeP is the currently active time synchronization mode.

4 **show sntp** again displays the SNTP configuration and shows that SNTP is now the currently active time synchronization mode and is configured for broadcast operation.

Figure 8-4. Example of Enabling SNTP Operation in Broadcast Mode

Enabling SNTP in Unicast Mode. Like broadcast mode, configuring SNTP for unicast mode enables SNTP. However, for Unicast operation, you must also specify the IP address of at least one SNTP server. The switch allows up to three unicast servers. You can use the Menu interface or the CLI to configure one server or to replace an existing Unicast server with another. To add a second or third server, you must use the CLI. For more on SNTP operation with multiple servers, see “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 22.

Syntax: timesync sntp — Selects SNTP as the time synchronization method.

sntp unicast — Configures the SNTP mode for Unicast operation.

sntp server <*ip-addr*> [*version*] — Specifies the SNTP server. The default server version is **3**.

no sntp server <*ip-addr*> Deletes the specified SNTP server.

Note
Deleting an SNTP server when only one is configured disables SNTP unicast operation.

For example, to select SNTP and configure it with unicast mode and an SNTP server at 10.28.227.141 with the default server version (3) and default poll interval (720 seconds):

```
HPswitch(config)# timesync sntpSelects SNTP.
HPswitch(config)# sntp unicastActivates SNTP in Unicast mode.
HPswitch(config)# sntp server 10.28.227.141Specifies the
                              SNTP server and accepts the
                              current SNTP server version
                              (default: 3).
```

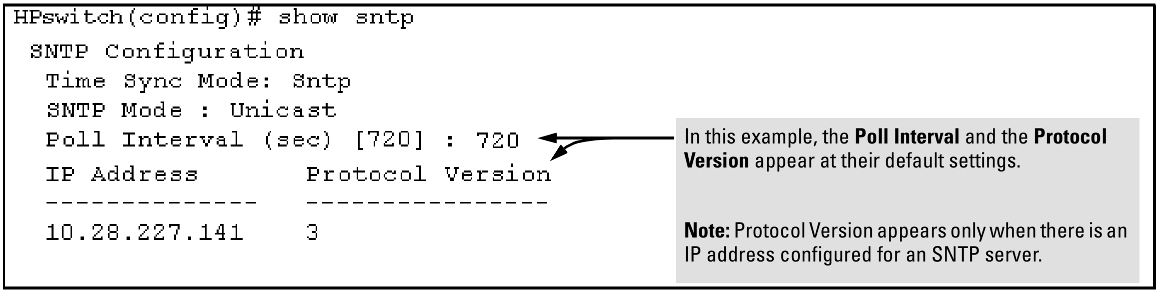

```
HPswitch(config)# show sntp

 SNTP Configuration

  Time Sync Mode: Sntp
  SNTP Mode : Unicast
  Poll Interval (sec) [720] : 720

  IP Address       Protocol Version
  --------------   ----------------
  10.28.227.141    3
```

In this example, the **Poll Interval** and the **Protocol Version** appear at their default settings.

Note: Protocol Version appears only when there is an IP address configured for an SNTP server.

Figure 8-5. Example of Configuring SNTP for Unicast Operation

If the SNTP server you specify uses SNTP version 4 or later, use the sntp server command to specify the correct version number. For example, suppose you learned that SNTP version 4 was in use on the server you specified above (IP address 10.28.227.141). You would use the following commands to delete the server IP address and then re-enter it with the correct version number for that server:

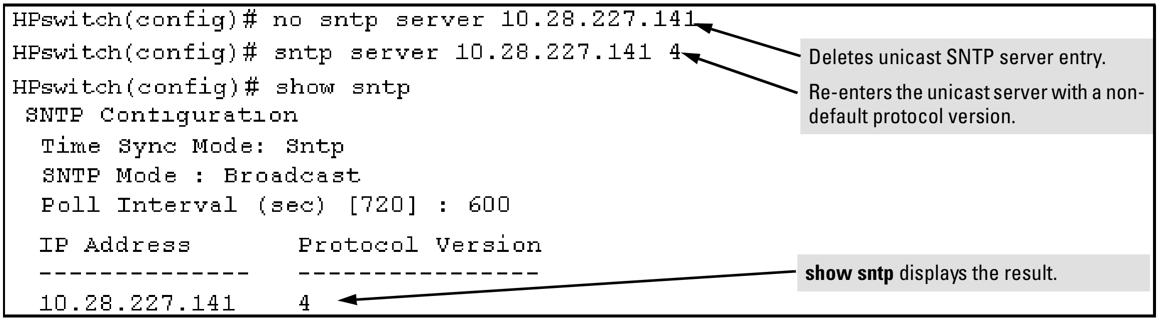

```
HPswitch(config)# no sntp server 10.28.227.141
HPswitch(config)# sntp server 10.28.227.141 4
HPswitch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 600

  IP Address       Protocol Version
  --------------   ----------------
  10.28.227.141    4
```

Deletes unicast SNTP server entry.

Re-enters the unicast server with a non-default protocol version.

show sntp displays the result.

Figure 8-6. Example of Specifying the SNTP Protocol Version Number

Changing the SNTP Poll Interval. This command lets you specify how long the switch waits between time polling intervals. The default is 720 seconds and the range is 30 to 720 seconds. (This parameter is separate from the poll interval parameter used for Timep operation.)

Syntax: sntp poll-interval <30 . . 720>

For example, to change the poll interval to 300 seconds:

```
HPswitch(config)# sntp poll-interval 300
```

Disabling Time Synchronization Without Changing the SNTP Configuration. The recommended method for disabling time synchronization is to use the **timesync** command. This halts time synchronization without changing your SNTP configuration.

Syntax: no timesync

For example, suppose SNTP is running as the switch's time synchronization protocol, with **Broadcast** as the SNTP mode and the factory-default polling interval. You would halt time synchronization with this command:

```
HPswitch(config)# no timesync
```

If you then viewed the SNTP configuration, you would see the following:

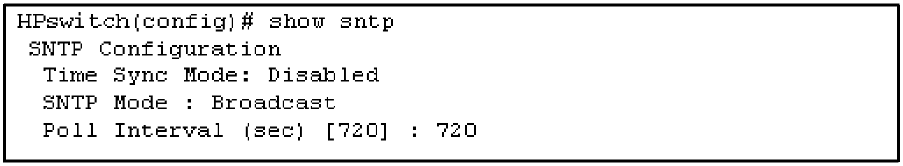

```
HPswitch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Disabled
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 720
```

Figure 8-7. Example of SNTP with Time Sychronization Disabled

Disabling the SNTP Mode. If you want to prevent SNTP from being used even if selected by **timesync** (or the Menu interface's **Time Sync Method** parameter), configure the SNTP mode as disabled.

Syntax: no sntp — Disables SNTP by changing the SNTP mode configuration to **Disabled**.

For example, if the switch is running SNTP in Unicast mode with an SNTP server at 10.28.227.141 and a server version of 3 (the default), **no sntp** changes the SNTP configuration as shown below, and disables time synchronization on the switch.

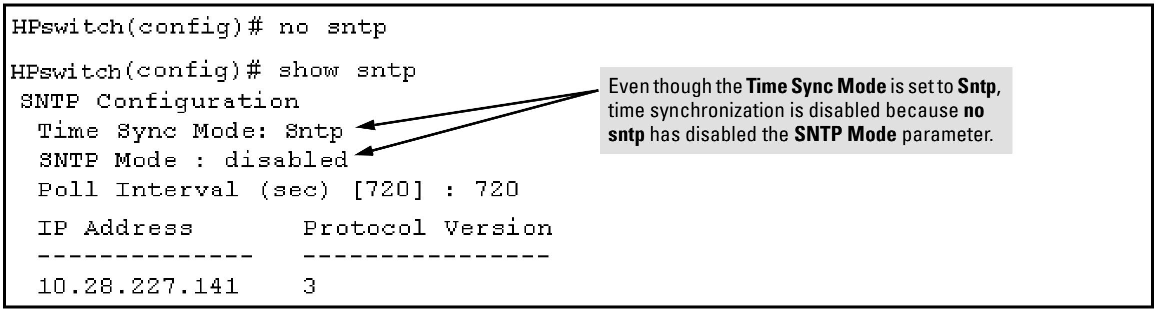

```
HPswitch(config)# no sntp

HPswitch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720

  IP Address       Protocol Version
  --------------   ----------------
  10.28.227.141    3
```

Even though the **Time Sync Mode** is set to **Sntp**, time synchronization is disabled because **no sntp** has disabled the **SNTP Mode** parameter.

Figure 8-8. Example of Disabling Time Synchronization by Disabling the SNTP Mode

TimeP: Viewing, Selecting, and Configuring

TimeP Feature	Default	Menu	CLI	Web
view the Timep time synchronization configuration	n/a	page 8-15	page 8-17	—
select Timep as the time synchronization method	TIMEP	page 8-13	pages 8-18 ff.	—
disable time synchronization	timep	page 8-15	page 8-20	—
enable the Timep mode	Disabled			—
DHCP	—	page 8-15	page 8-18	—
manual	—	page 8-16	page 8-19	—
none/disabled	—	page 8-15	page 8-21	—
change the SNTP poll interval	720 minutes	page 8-16	page 8-20	—

Table 8-2.Timep Parameters

SNTP Parameter	Operation
Time Sync Method	Used to select either TIMEP (the default), SNTP, or None as the time synchronization method.
Timep Mode	
Disabled	The Default. Timep does not operate, even if specified by the Menu interface **Time Sync Method** parameter or the CLI **timesync** command.
DHCP	When Timep is selected as the time synchronization method, the switch attempts to acquire a Timep server IP address via DHCP. If the switch receives a server address, it polls the server for updates according to the Timep poll interval. If the switch does not receive a Timep server IP address, it cannot perform time synchronization updates.
Manual	When Timep is selected as the time synchronization method, the switch attempts to poll the specified server for updates according to the Timep poll interval. If the switch fails to receive updates from the server, time synchronization updates do not occur.
Server Address	Used only when the **TimeP Mode** is set to **Manual**. Specifies the IP address of the TimeP server that the switch accesses for time synchronization updates. You can configure one server.
Poll Interval (minutes)	Default: 720 minutes. Specifies the interval the switch waits between attempts to poll the TimeP server for updates.

Menu: Viewing and Configuring TimeP

To View, Enable, and Modify the TimeP Protocol:

1. From the Main Menu, select:

   2. Switch Configuration...

   1. System Information

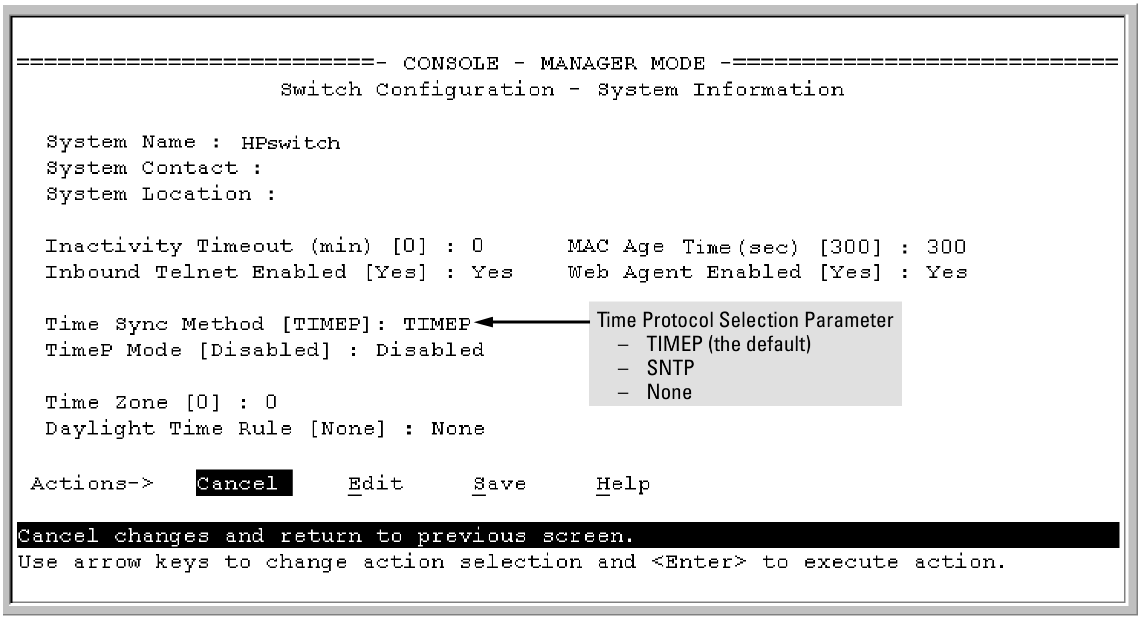

```
=========================- CONSOLE - MANAGER MODE -============================
                   Switch Configuration - System Information

  System Name : HPswitch
  System Contact :
  System Location :

  Inactivity Timeout (min) [0] : 0      MAC Age Time(sec) [300] : 300
  Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes

  Time Sync Method [TIMEP]: TIMEP
  TimeP Mode [Disabled] : Disabled

  Time Zone [0] : 0
  Daylight Time Rule [None] : None

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Time Protocol Selection Parameter
– TIMEP (the default)
– SNTP
– None

Figure 8-9. The System Information Screen (Default Values)

2. Press [E] (for **Edit**). The cursor moves to the **System Name** field.

3. Use [↓] to move the cursor to the **Time Sync Method** field.

4. If **TIMEP** is not already selected, use the Space bar to select **TIMEP**, then press [↓] once to display and move to the **TimeP Mode** field.

5. Do one of the following:

   - Use the Space bar to select the **DHCP** mode, then press [↓] to move the cursor to the **Poll Interval** field, and go to step 6.

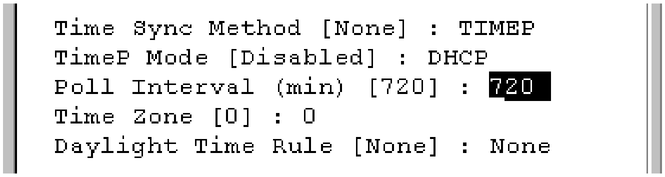

```
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : DHCP
Poll Interval (min) [720] : 720
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

- Use the Space bar to select the **Manual** mode.

  i. Press ⇥ to move the cursor to the **Server Address** field.

  ii. Enter the IP address of the TimeP server you want the switch to use for time synchronization.

  Note: This step replaces any previously configured TimeP server IP address.

  iii. Press ⇥ to move the cursor to the **Poll Interval** field, then go to step 6.

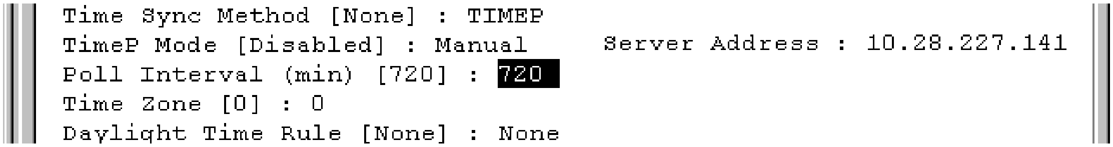

```
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Manual     Server Address : 10.28.227.141
Poll Interval (min) [720] : 720
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

6. In the **Poll Interval** field, enter the time in minutes that you want for a TimeP Poll Interval.

Press [Enter] to return to the Actions line, then [S] (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

CLI: Viewing and Configuring TimeP

CLI Commands Described in this Section

Command	Page
show timep	8-17
[no] timesync	8-18 ff., 8-20
ip timep	
dhcp	8-18
manual	8-19
server <*ip-addr*>	8-19
interval	8-20
no ip timep	8-21

This section describes how to use the CLI to view, enable, and configure TimeP parameters.

Viewing the Current TimeP Configuration

This command lists both the time synchronization method (TimeP, SNTP, or None) and the TimeP configuration, even if SNTP is not the selected time protocol.

Syntax: show timep

For example, if you configure the switch with TimeP as the time synchronization method, then enable TimeP in DHCP mode with the default poll interval, **show timep** lists the following:

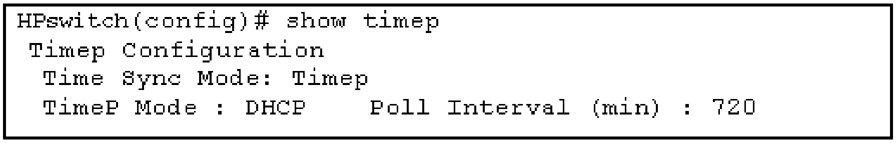

```
HPswitch(config)# show timep
 Timep Configuration
  Time Sync Mode: Timep
  TimeP Mode : DHCP    Poll Interval (min) : 720
```

Figure 8-10. Example of TimeP Configuration When TimeP Is the Selected Time Synchronization Method

If SNTP is the selected time synchronization method), **show timep** still lists the TimeP configuration even though it is not currently in use:

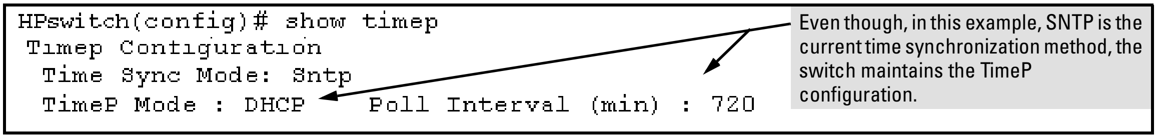

```
HPswitch(config)# show timep
 Timep Configuration
  Time Sync Mode: Sntp
  TimeP Mode : DHCP     Poll Interval (min) : 720
```

Even though, in this example, SNTP is the current time synchronization method, the switch maintains the TimeP configuration.

Figure 8-11. Example of SNTP Configuration When SNTP Is Not the Selected Time Synchronization Method

Configuring (Enabling or Disabling) the TimeP Mode

Enabling the TimeP mode means to configure it for either broadcast or unicast mode. Remember that to run TimeP as the switch's time synchronization protocol, you must also select TimeP as the time synchronization method by using the CLI timesync command (or the Menu interface **Time Sync Method** parameter).

Syntax:	timesync timep	Selects TimeP as the time protocol.
	ip timep < dhcp | manual >	Enables the selected TimeP mode.
	no ip timep	Disables the TimeP mode.
	no timesync	Disables the time protocol.

Enabling TimeP in DHCP Mode. Because the switch provides a TimeP polling interval (default: 720 minutes), you need only these two commands for a minimal TimeP DHCP configuration:

Syntax:	timesync timep	Selects TimeP as the time synchronization method.
	ip timep dhcp	Configures DHCP as the TimeP mode.

For example, suppose:

- Time synchronization is configured for SNTP.
- You want to:
  1. View the current time synchronization.
  2. Select TimeP as the time synchronization mode.
  3. Enable TimeP for DHCP mode.
  4. View the TimeP configuration.

The commands and output would appear as follows:

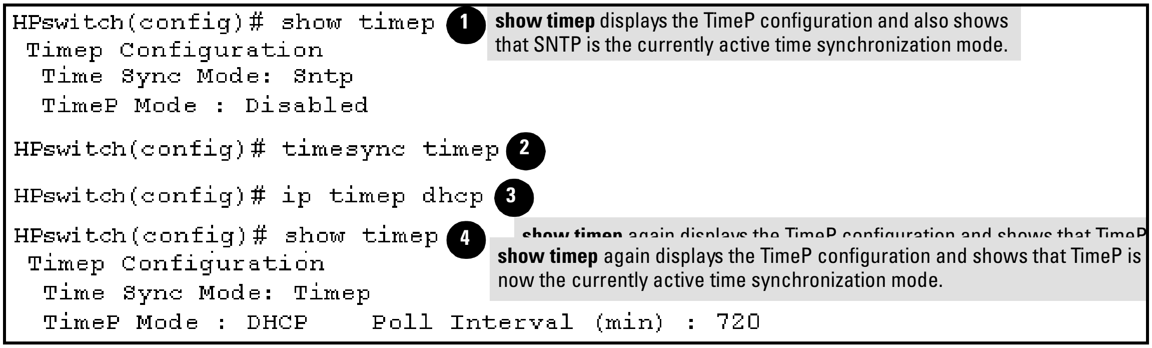

```
HPswitch(config)# show timep 1
 Timep Configuration
  Time Sync Mode: Sntp
  TimeP Mode : Disabled

HPswitch(config)# timesync timep 2

HPswitch(config)# ip timep dhcp 3

HPswitch(config)# show timep 4
 Timep Configuration
  Time Sync Mode: Timep
  TimeP Mode : DHCP    Poll Interval (min) : 720
```

show timep displays the TimeP configuration and also shows that SNTP is the currently active time synchronization mode.

show timep again displays the TimeP configuration and shows that TimeP is now the currently active time synchronization mode.

Figure 8-12. Example of Enabling TimeP Operation in DHCP Mode

Enabling Timep in Manual Mode. Like DHCP mode, configuring TimeP for **Manual** mode enables TimeP. However, for manual operation, you must also specify the IP address of the TimeP server. (The switch allows only one TimeP server.) To enable the TimeP protocol:

Syntax:	timesync timep	Selects Timep.
	ip timep manual <*ip-addr*>	Activates TimeP in Manual mode with a specified TimeP server.
	no ip timep	Disables TimeP.

Note

To change from one TimeP server to another, you must (1) use the **no ip timep** command to disable TimeP mode, and then reconfigure TimeP in Manual mode with the new server IP address.

For example, to select TimeP and configure it for manual operation using a TimeP server address of 10.28.227.141 and the default poll interval (720 minutes, assuming the TimeP poll interval is already set to the default):

```
HPswitch(config)# timesync timep     Selects TimeP.
HPswitch(config)# ip timep manual    Activates TimeP in Manual
   10.28.227.141                     mode.
```

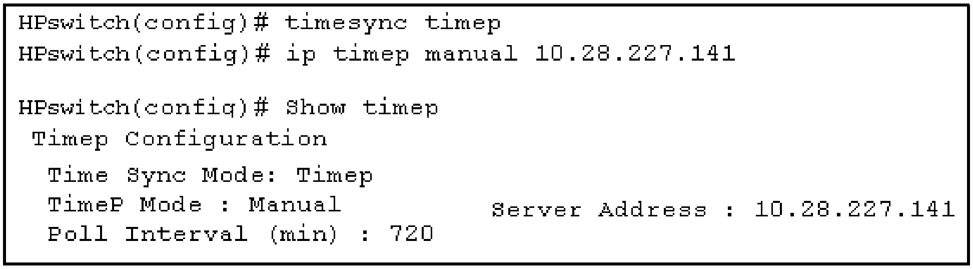

```
HPswitch(config)# timesync timep
HPswitch(config)# ip timep manual 10.28.227.141

HPswitch(config)# Show timep
 Timep Configuration

  Time Sync Mode: Timep
  TimeP Mode : Manual              Server Address : 10.28.227.141
  Poll Interval (min) : 720
```

Figure 8-13. Example of Configuring Timep for Manual Operation

Changing the TimeP Poll Interval. This command lets you specify how long the switch waits between time polling intervals. The default is 720 minutes and the range is 1 to 9999 minutes. (This parameter is separate from the poll interval parameter used for SNTP operation.)

Syntax: ip timep dhcp interval <1 . . 9999>
ip timep manual interval <1 . . 9999>

For example, to change the poll interval to 60 minutes:

```
HPswitch(config)# ip timep interval 60
```

Disabling Time Synchronization Without Changing the TimeP Configuration. The recommended method for disabling time synchronization is to use the **timesync** command. This halts time synchronization without changing your TimeP configuration.

Syntax: no timesync

For example, suppose TimeP is running as the switch's time synchronization protocol, with **DHCP** as the TimeP mode, and the factory-default polling interval. You would halt time synchronization with this command:

```
HPswitch(config)# no timesync
```

If you then viewed the TimeP configuration, you would see the following:

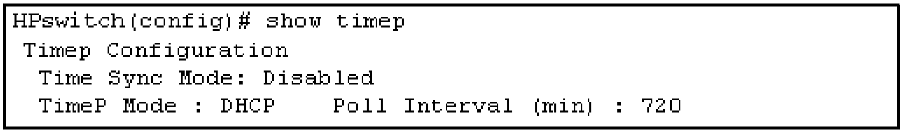

```
HPswitch(config)# show timep
 Timep Configuration
  Time Sync Mode: Disabled
  TimeP Mode : DHCP    Poll Interval (min) : 720
```

Figure 8-14. Example of TimeP with Time Sychronization Disabled

Disabling the TimeP Mode. Disabling the TimeP mode means to configure it as disabled. (Disabling TimeP prevents the switch from using it as the time synchronization protocol, even if it is the selected **Time Sync Method** option.)

Syntax: no ip timep   Disables TimeP by changing the TimeP mode configuration to **Disabled**.

For example, if the switch is running TimeP in DHCP mode, **no ip timep** changes the TimeP configuration as shown below, and disables time synchronization on the switch.

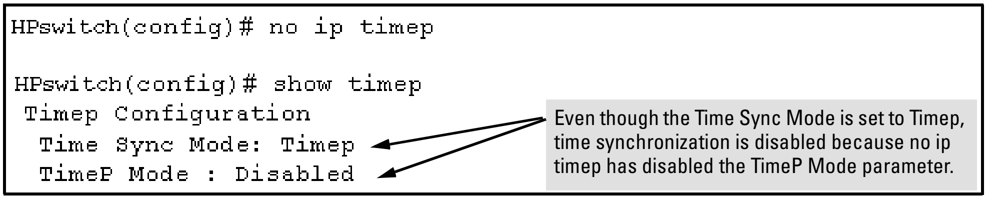

```
HPswitch(config)# no ip timep

HPswitch(config)# show timep
 Timep Configuration
  Time Sync Mode: Timep
  TimeP Mode : Disabled
```

Even though the Time Sync Mode is set to Timep, time synchronization is disabled because no ip timep has disabled the TimeP Mode parameter.

Figure 8-15. Example of Disabling Time Synchronization by Disabling the TimeP Mode Parameter

SNTP Unicast Time Polling with Multiple SNTP Servers

When running SNTP unicast time polling as the time synchronization method, the switch requests a time update from the server you configured with either the Server Address parameter in the menu interface, or the primary server in a list of up to three SNTP servers configured using the CLI. If the switch does not receive a response from the primary server after three consecutive polling intervals, the switch tries the next server (if any) in the list. If the switch tries all servers in the list without success, it sends an error message to the Event Log and reschedules to try the address list again after the configured **Poll Interval** time has expired.

Address Prioritization

If you use the CLI to configure multiple SNTP servers, the switch prioritizes them according to the decimal values of their IP addresses. That is, the switch compares the decimal value of the octets in the addresses and orders them accordingly, with the lowest decimal value assigned as the primary address, the second-lowest decimal value assigned as the next address, and the third-lowest decimal value as the last address. If the first octet is the same between two of the addresses, the second octet is compared, and so on. For example:

SNTP Server IP Address	Server Ranking According to Decimal Value of IP Address
10.28.227.141	Primary
10.28.227.153	Secondary
10.29.227.100	Tertiary

Adding and Deleting SNTP Server Addresses

Adding Addresses. As mentioned earlier, you can configure one SNTP server address using either the Menu interface or the CLI. To configure a second and third address, you must use the CLI. For example, suppose you have already configured the primary address in the above table (10.28.227.141). To configure the remaining two addresses, you would do the following:

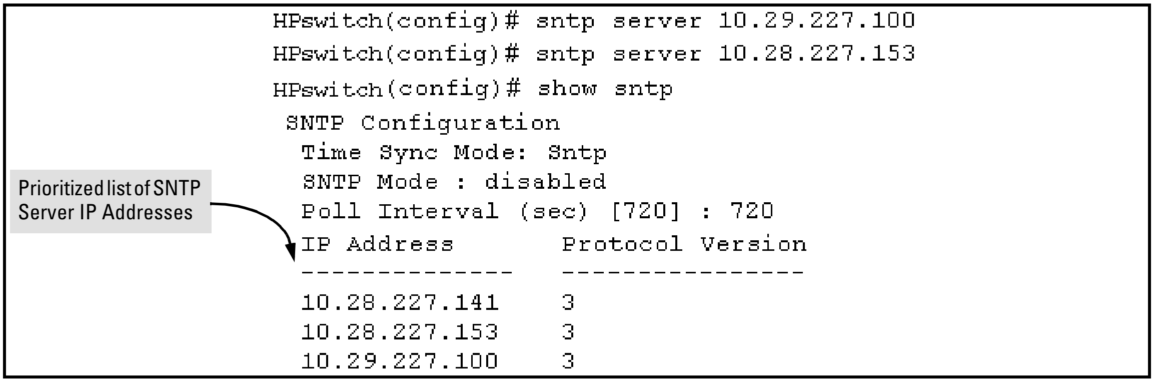

```
HPswitch(config)# sntp server 10.29.227.100
HPswitch(config)# sntp server 10.28.227.153
HPswitch(config)# show sntp
 SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720
  IP Address       Protocol Version
  --------------   ----------------
  10.28.227.141    3
  10.28.227.153    3
  10.29.227.100    3
```

Prioritized list of SNTP Server IP Addresses

Figure 8-16. Example of SNTP Server Address Prioritization

Note

If there are already three SNTP server addresses configured on the switch, and you want to use the CLI to replace one of the existing addresses with a new one, you must delete the unwanted address before you configure the new one.

Deleting Addresses. To delete an address, you must use the CLI. If there are multiple addresses and you delete one of them, the switch re-orders the address priority. (See “Address Prioritization” on page 22.)

Syntax: no sntp server <*ip-addr*>

For example, to delete the primary address in the above example (and automatically convert the secondary address to primary):

```
HPswitch(config)# no sntp server 10.28.227.141
```

Menu: Operation with Multiple SNTP Server Addresses Configured

When you use the Menu interface to configure an SNTP server IP address, the new address writes over the current primary address, if one is configured. If there are multiple addresses configured, the switch re-orders the addresses according to the criteria described under "Address Prioritization" on page 22. For example, suppose the switch already has the following three SNTP server IP addresses configured.

- 10.28.227.141 (primary)
- 10.28.227.153 (secondary)
- 10.29.227.100 (tertiary)

If you use the Menu interface to add 10.28.227.160, the new prioritized list will be:

New Address List	Address Status
10.28.227.153	New Primary (The former primary, 10.28.227.141 was deleted when you used the menu to add 10.28.227.160.)
10.28.227.160	New Secondary
10.29.227.100	Same Tertiary (This address still has the highest decimal value.)

SNTP Messages in the Event Log

If an SNTP time change of more than three seconds occurs, the switch's event log records the change. SNTP time changes of less than three seconds do not appear in the Event Log.

Optimizing Traffic Flow with Port Controls, Port Trunking, and Filters

Contents

Overview

This chapter includes:

- Configuring ports to non-default settings (page 9-2)

  These settings include enable/disable, mode (speed and duplex), flow control, port-trunk group, and port-trunk type. You can also set a broadcast limit that applies to all ports on the switch.

- Port aggregation: Creating and modifying a port trunk group (page 9-10)

  You can configure static and dynamic trunks. Includes non-protocol trunks, LACP (802.3ad) trunks, and FEC trunks.

- Traffic/Security Filters: Configuring filters based on source-port, protocol, and multicast address.

Viewing Port Status and Configuring Port Parameters

Port Status and Configuration Features

Feature	Default	Menu	CLI	Web
viewing port status	n/a	page 9-5	page 9-6	page 9-9
configuring ports	See Table 9-1 on pages 9-3 and 9-4	page 9-5	page 9-8	page 9-9

Note On Connecting Transceivers to Fixed-Configuration Devices

If the switch either fails to show a link between an installed transceiver and another device, or demonstrates errors or other unexpected behavior on the link, check the port configuration on both devices for a speed and/or duplex (mode) mismatch. To check the mode setting for a port on a Series 5300XL Switches, use either the Port Status screen in the menu interface (page 9-5) or **show interfaces brief** in the CLI (page 9-6).

Table 9-1. Status and Parameters for Each Port Type

Status or Parameter	Description
Enabled	**Yes** (default): The port is ready for a network connection. **No:** The port will not operate, even if properly connected in a network. Use this setting, for example, if the port needs to be shut down for diagnostic purposes or while you are making topology changes.
Status (read-only)	**Up**: The port senses a link beat. **Down**: The port is not enabled, has no cables connected, or is experiencing a network error. For troubleshooting information, see the installation manual you received with the switch. See also appendix C, “Troubleshooting” (in this manual).
Mode	The port’s speed and duplex (data transfer operation) setting. 10/100Base-T ports: • Auto (default): Senses speed and negotiates with the port at the other end of the link for data transfer operation (half-duplex or full-duplex). **Note:** Ensure that the device attached to the port is configured for the same setting that you select here. Also, if “Auto” is used, the device to which the port is connected must operate in compliance with the IEEE 802.3u “Auto Negotiation” standard for 100Base-T networks. If the other device does not comply with the 802.3u standard, or is not set to Auto, then the port configuration on the switch must be manually set to match the port configuration on the other device. To see what the switch negotiates for the Auto setting, use the CLI **show interfaces** command or the “**3. Port Status”** option under “**1. Status and Counters”** in the menu interface. • Auto-10: Allows the port to negotiate between half-duplex (HDx) and full-duplex (FDx) while keeping speed at 10 Mbps. Also negotiates flow control (enabled or disabled). HP recommends Auto-10 for links between 10/100 auto-sensing ports connected with Cat 3 cabling. (Cat 5 cabling is required for 100 Mbps links.). • 10HDx:10 Mbps, Half-Duplex • 10FDx: 10 Mbps, Full-Duplex • 100HDx: 100 Mbps, Half-Duplex • 100FDx: 100 Mbps, Full-Duplex
	100FX ports: • 100HDx: 100 Mbps, Half-Duplex • 100FDx (default): 100 Mbps, Full-Duplex

Status or Parameter	Description
	100/1000Base-T ports: • Auto (default): Senses speed and negotiates with the port at the other end of the link for port operation (MDI-X or MDI). To see what the switch negotiates for the Auto setting, use the CLI **show interfaces brief** command or the "**3. Port Status"** option under "**1. Status and Counters"** in the menu interface. • Auto-100: Uses 100 Mbps and negotiates with the port at the other end of the link for other port operation features. • Auto-1000: Uses 1000 Mbps and negotiates with the port at the other end of the link for other port operation features. • 100Hdx: Uses 100 Mbps, half-duplex. • 100Fdx: Uses 100 Mbps, Full-Duplex **Notes:** • Changing the port speed on a transceiver port requires a reboot of the switch. • Ensure that the device attached to the port is configured for the same setting that you select here. Also, if "Auto" is used, the device to which the port is connected must also be configured to "Auto" and operate in compliance with the IEEE 802.3ab "Auto Negotiation" standard for 1000Base-T networks.
	Gigabit fiber-optic ports (Gigabit-SX and Gigabit-LX): • 1000FDx: 1000 Mbps (1 Gbps), Full Duplex only • Auto (default): The port operates at 1000FDx and auto-negotiates flow control with the device connected to the port.
Flow Control	• Disabled (default): The port will not generate flow control packets, and drops any flow control packets it receives. • Enabled: The port uses 802.3x Link Layer Flow Control, generates flow control packets, and processes received flow control packets. With the port mode set to Auto (the default) and Flow Control enabled, the switch negotiates Flow Control on the indicated port. If the port mode is not set to Auto, or if Flow Control is disabled on the port, then Flow Control is not used.
Group (menu) or Trunk Group (CLI)	Menu Interface: Specifies the static trunk group, if any, to which a port belongs. CLI: Appears in the **show lacp** command output to show the LACP trunk, if any, to which a port belongs. **Note:** An LACP trunk requires a full-duplex link. In most cases, HP recommends that you leave the port Mode setting at Auto (the default). See the LACP Note on page 9-12. *For more on port trunking, see "Port Trunking" on page 9-10.*
Type	This parameter appears in the CLI **show trunk** listing and, for a port in a trunk group, specifies the type of trunk group. The default Type is passive LACP, which can be displayed by using the CLI **show lacp** command. *For more on port trunking, see "Port Trunking" on page 9-10.*
Broadcast Limit	Reduces the bandwidth available for broadcast and multicast traffic on all ports on the switch. Any broadcast or multicast traffic overload will be dropped. This feature is not appropriate for networks that require high levels of IPX or RIP broadcast traffic.

Menu: Viewing Port Status and Configuring Port Parameters

From the menu interface, you can configure and view all port parameter settings and view all port status indicators.

Using the Menu To View Port Status. The menu interface displays the status for ports and (if configured) a trunk group.

From the Main Menu, select:

1. Status and Counters . . .

4. Port Status

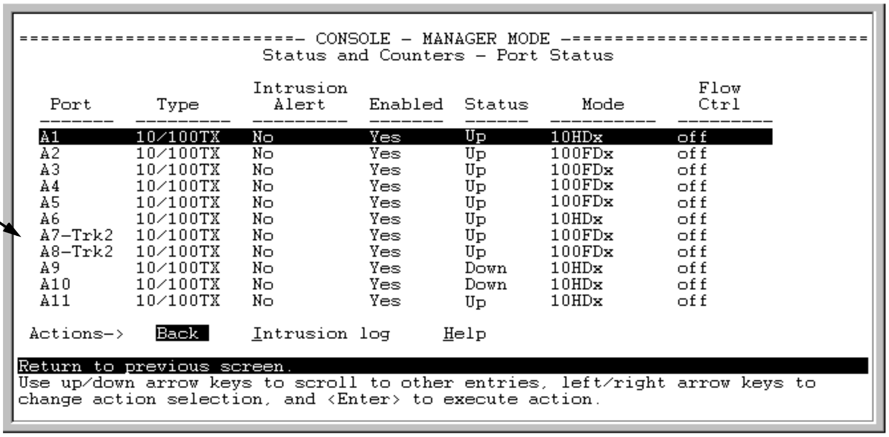

```
=========================- CONSOLE - MANAGER MODE -===========================
                     Status and Counters - Port Status

                           Intrusion                                 Flow
  Port       Type           Alert     Enabled  Status     Mode        Ctrl
  -------    ---------    ---------  -------  ------   ----------   ---------
  A1         10/100TX     No          Yes      Up       10HDx        off
  A2         10/100TX     No          Yes      Up       100FDx       off
  A3         10/100TX     No          Yes      Up       100FDx       off
  A4         10/100TX     No          Yes      Up       100FDx       off
  A5         10/100TX     No          Yes      Up       100FDx       off
  A6         10/100TX     No          Yes      Up       10HDx        off
  A7-Trk2    10/100TX     No          Yes      Up       100FDx       off
  A8-Trk2    10/100TX     No          Yes      Up       100FDx       off
  A9         10/100TX     No          Yes      Down     10HDx        off
  A10        10/100TX     No          Yes      Down     10HDx        off
  A11        10/100TX     No          Yes      Up       10HDx        off

 Actions->   Back      Intrusion log     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

In this example, ports A7 and A8 have previously been configured as a trunk group.

Figure 9-1. Example of the Port Status Screen

Using the Menu To Configure Ports.

Note

The menu interface uses the same screen for configuring both individual ports and port trunk groups. For information on port trunk groups, see "Port Trunking" on page 9-10.

1. From the Main Menu, Select:

   2. Switch Configuration...

   2. Port/Trunk Settings

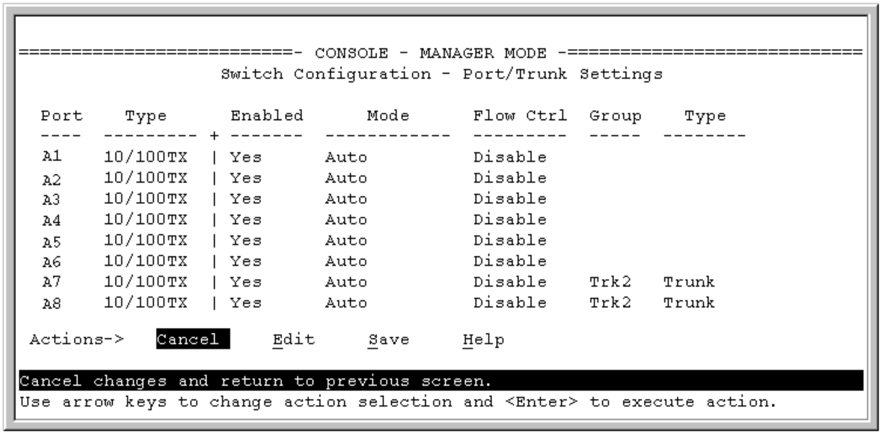

```
=========================- CONSOLE - MANAGER MODE -===========================
                  Switch Configuration - Port/Trunk Settings

  Port    Type      Enabled      Mode       Flow Ctrl  Group    Type
  ----  --------- + -------  ------------  ---------  -----  --------
  A1    10/100TX  | Yes      Auto          Disable
  A2    10/100TX  | Yes      Auto          Disable
  A3    10/100TX  | Yes      Auto          Disable
  A4    10/100TX  | Yes      Auto          Disable
  A5    10/100TX  | Yes      Auto          Disable
  A6    10/100TX  | Yes      Auto          Disable
  A7    10/100TX  | Yes      Auto          Disable    Trk2   Trunk
  A8    10/100TX  | Yes      Auto          Disable    Trk2   Trunk

 Actions->  Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 9-2. Example of Port/Trunk Settings with a Trunk Group Configured

2. Press [E] (for Edit). The cursor moves to the **Enabled** field for the first port.
3. Refer to the online help provided with this screen for further information on configuration options for these features.
4. When you have finished making changes to the above parameters, press [Enter], then press [S] (for **Save**).

CLI: Viewing Port Status and Configuring Port Parameters

Port Status and Configuration Commands

show interfaces brief	below
show interface config	page 9-7
interface	page 9-8

From the CLI, you can configure and view all port parameter settings and view all port status indicators.

Using the CLI To View Port Status. Use the following commands to display port status and configuration:

- **show interfaces brief**: Lists the full status and configuration for all ports on the switch.
- **show interface config**: Lists a subset of the data shown by the **show interfaces** command (above); that is, only the enabled/disabled, mode, and flow control status for all ports on the switch.

Syntax: show interfaces brief
show interface config

The next two figures list examples of the output of the above two commands for the same port configuration.

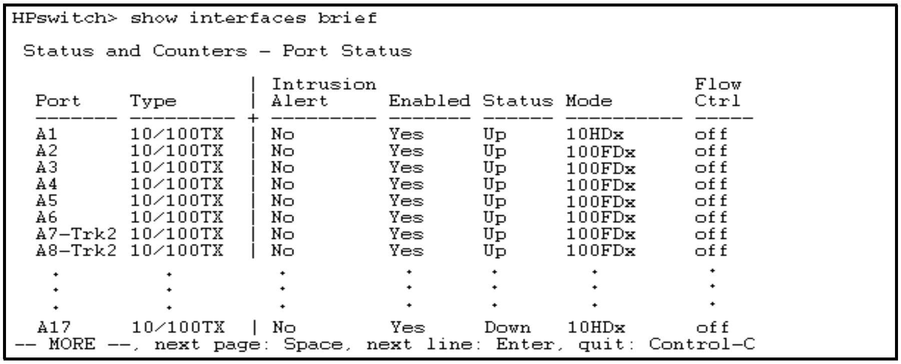

```
HPswitch> show interfaces brief

 Status and Counters - Port Status

                     | Intrusion                                  Flow
  Port     Type      | Alert     Enabled Status Mode        Ctrl
  -------  --------- + --------- ------- ------ ----------  -----
  A1       10/100TX  | No        Yes     Up     10HDx       off
  A2       10/100TX  | No        Yes     Up     100FDx      off
  A3       10/100TX  | No        Yes     Up     100FDx      off
  A4       10/100TX  | No        Yes     Up     100FDx      off
  A5       10/100TX  | No        Yes     Up     100FDx      off
  A6       10/100TX  | No        Yes     Up     100FDx      off
  A7-Trk2  10/100TX  | No        Yes     Up     100FDx      off
  A8-Trk2  10/100TX  | No        Yes     Up     100FDx      off
   .         .         .          .      .       .           .
   .         .         .          .      .       .           .
   .         .         .          .      .       .           .
  A17      10/100TX  | No        Yes     Down   10HDx       off
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

Figure 9-3. Example of a Show Interface Command Listing

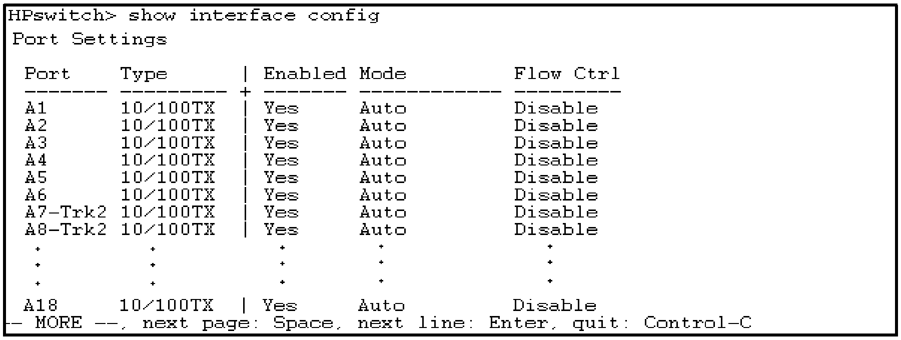

```
HPswitch> show interface config
 Port Settings

  Port     Type      | Enabled Mode         Flow Ctrl
  -------  --------- + ------- ------------ ---------
  A1       10/100TX  | Yes     Auto         Disable
  A2       10/100TX  | Yes     Auto         Disable
  A3       10/100TX  | Yes     Auto         Disable
  A4       10/100TX  | Yes     Auto         Disable
  A5       10/100TX  | Yes     Auto         Disable
  A6       10/100TX  | Yes     Auto         Disable
  A7-Trk2  10/100TX  | Yes     Auto         Disable
  A8-Trk2  10/100TX  | Yes     Auto         Disable
   .         .          .       .              .
   .         .          .       .              .
   .         .          .       .              .
  A18      10/100TX  | Yes     Auto         Disable
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

Figure 9-4. Example of a Show Interface Config Command Listing

Using the CLI To Configure Ports. You can configure one or more of the following port parameters. For details on each option, see Table 9-1 on page 9-3.

Syntax: [no] interface <[ethernet] *port-list*>

```
[no] interface <[ethernet] port-list>
     disable | enable
     speed-duplex
          <auto-10 |10-full | 10-half | 100-full | 100-half |auto|1000-full |>
     flow-control
```

Note that in the above syntax you can substitute an "**int**" for "**interface**" and an "**e**" for "**ethernet**"; that is **int e <*port-list*>**.

For example, to configure ports C1 through C3 and port C6 for 100Mbps full-duplex, you would enter these commands:

```
HPswitch(config)# int e c1-c3,c6 speed-duplex 100-full
```

Similarly, to configure a single port with the settings in the above command, you could either enter the same command with only the one port identified, or go to the *context level* for that port and then enter the command. For example, to enter the context level for port C6 and then configure that port for 100FDx:

```
HPswitch(config)# int e c6
HPswitch(eth-C6)# speed-duplex 100-full
```

If port C8 was disabled, and you wanted to enable it and configure it for 100FDx with flow-control active, you could do so with either of the following command sets.

- These commands enable and configure port C8 from the config level:
  ```
  HPswitch(config)# int e c8 enable
  HPswitch(config)# int e c8 speed-duplex 100-full
  HPswitch(config)# int e c8 flow-control
  ```
- These commands select the context level for port C8 and then apply all of the configuration commands to port C8:
  ```
  HPswitch(config)# int e c8
  HPswitch(eth-C8)# enable
  HPswitch(eth-C8)# speed-duplex 100-full
  HPswitch(eth-C8)# flow-control
  ```

Configuring a Broadcast Limiting on the Switch. Executing this command configures broadcast limiting for all ports on the switch.

Syntax: broadcast-limit

```
HPswitch(config)# broadcast-limit
```

To display the current broadcast limit setting, use one of the following commands:

`HPswitch# show config` Displays the startup-config file. The broadcast limit setting appears here if enabled and saved to the startup-config file.

`HPswitch# show running-config` Displays the running-config file. The broadcast limit setting appears here if enabled. If the setting is not also saved to the startup-config file, rebooting the switch returns broad cast limit to the setting currently in the startup-config file.

Web: Viewing Port Status and Configuring Port Parameters

In the web browser interface:

1. Click on the **Configuration** tab.
2. Click on [Port Configuration].
3. Select the ports you want to modify and click on [Modify Selected Ports].
4. After you make the desired changes, click on [Apply Settings].

Note that the web browser interface displays an existing port trunk group. However, to configure a port trunk group, you must use the CLI or the menu interface. For more on this topic, see "Port Trunking" on page 9-10.

Port Trunking

Port Status and Configuration Features

Feature	Default	Menu	CLI	Web
viewing port trunks	n/a	page 9-16	page 9-18	page 9-24
configuring a static trunk group	none	page 9-16	page 9-22	—
configuring a dynamic LACP trunk group	LACP passive	—	page 9-23	—

Port trunking allows you to assign up to four physical links to one logical link (trunk) that functions as a single, higher-speed link providing dramatically increased bandwidth. This capability applies to connections between backbone devices as well as to connections in other network areas where traffic bottlenecks exist. A *trunk group* is a set of up to four ports configured as members of the same port trunk. Note that the ports in a trunk group do not have to be consecutive. For example:

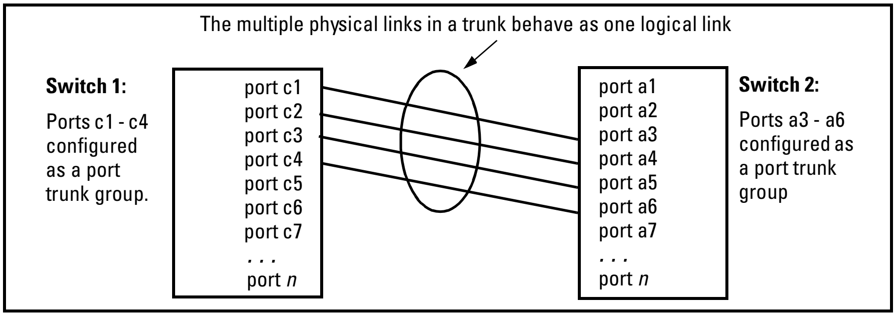
The multiple physical links in a trunk behave as one logical link
Switch 1:
Ports c1 - c4 configured as a port trunk group.
port c1
port c2
port c3
port c4
port c5
port c6
port c7
. . .
port *n*
port a1
port a2
port a3
port a4
port a5
port a6
port a7
. . .
port *n*
Switch 2:
Ports a3 - a6 configured as a port trunk group

Figure 9-5. Conceptual Example of Port Trunking

With full-duplex operation in a four-port trunk group, trunking enables the following bandwidth capabilities:

Table 9-2. Bandwidth Capacity for Trunk Groups Configured for Full-Duplex

	10 Mbps Links	100 Mbps Links	1000 Mbps Links
2 Ports	Up to 40 Mbps	Up to 400 Mbps	Up to 4000 Mbps
3 Ports	Up to 60 Mbps	Up to 600 Mbps	Up to 6000 Mbps
4 Ports	Up to 80 Mbps	Up to 800 Mbps	Up to 8000 Mbps

Port Connections and Configuration: All port trunk links must be point-to-point connections between the Series 5300XL Switches and another switch, router, server, or workstation configured for port trunking. No intervening, non-trunking devices are allowed. It is important to note that ports on both ends of a port trunk group must have the same mode (speed and duplex) and flow control settings.

Note

Link Connections. The switch does not support port trunking through an intermediate, non-trunking device such as a hub, or using more than one media type in a port trunk group. Similarly, all links in the same trunk group must have the same speed, duplex, and flow control.

Port Security Restriction. Port security does not operate on a trunk group. If you configure port security on one or more ports that are later added to a trunk group, the switch will reset the port security parameters for those ports to the factory-default configuration.

Caution

To avoid broadcast storms or loops in your network while configuring a trunk, first disable or disconnect all ports you want to add to or remove from the trunk. After you finish configuring the trunk, enable or re-connect the ports.

Series 5300XL Switches Port Trunk Features and Operation

The Series 5300XL Switches offers these options for port trunking:

- LACP (IEEE 802.3ad—page 9-25)
- Trunk (non-protocol—page 9-30)
- FEC (Fast EtherChannel®—page 9-30)

The switch supports 36 trunk groups of up to four ports each. (Using the Link Aggregation Control Protocol—LACP—option, you can include standby trunked ports in addition to the maximum of four actively trunking ports.)

LACP Note

LACP operation requires full-duplex (FDx) links. For most installations, HP recommends that you leave the port Mode settings at **Auto** (the default). LACP also operates with **Auto-10**, **Auto-100**, and **Auto-1000** (if negotiation selects FDx); **10FDx**, **100FDx**, and **1000FDx** settings.

Fault Tolerance: If a link in a port trunk fails, the switch redistributes traffic originally destined for that link to the remaining links in the trunk. The trunk remains operable as long as there is at least one link in operation. If a link is restored, that link is automatically included in the traffic distribution again. The LACP option also offers a standby link capability, which enables you to keep links in reserve for service if one or more of the original active links fails. See "Trunk Group Operation Using LACP" on page 9-25.)

Trunk Configuration Methods

Dynamic LACP Trunk: The switch automatically negotiates trunked links between LACP-configured ports on separate devices, and offers one dynamic trunk option: LACP. To configure the switch to initiate a dynamic LACP trunk with another device, use the **interface ethernet** command in the CLI to set the default LACP option to **Active** on the ports you want to use for the trunk. For example, the following command sets ports C1-C4 to LACP active:

```
HPswitch(config) int e c1-c4 lacp active
```

Note that the above example works if the ports are not already operating in a trunk. To change the LACP option on ports already operating as a trunk, you must first remove them from the trunk. For example, if ports C1 - C4 were LACP-active and operating in a trunk with another device, you would do the following to change them to LACP-passive:

```
HPswitch(config)# no int e c1-c4 lacp
```

Removes the ports from the trunk.

```
HPswitch(config)# int e c1-c4 lacp passive
```

Configures LACP passive.

Static Trunk: The switch uses the links you configure with the Port/Trunk Settings screen in the menu interface or the **trunk** command in the CLI to create a static port trunk. The switch offers three types of static trunks: LACP, Trunk, and FEC.

Table 9-3. Trunk Types Used in Static and Dynamic Trunk Groups

Trunking Method	LACP	Trunk	FEC
Dynamic	Yes	No	No
Static	Yes	Yes	Yes

Table 9-4. Trunk Configuration Protocols

Protocol	Trunking Options
LACP (802.3ad)	Provides dynamic and static LACP trunking options. • **Dynamic LACP** — Use the switch-negotiated dynamic LACP trunk when: – The port on the other end of the trunk link is configured for Active or Passive LACP. – You want to achieve fault-tolerance for high-availability applications where you want a four-link trunk with one or more standby links available in case an active link goes down. (Both ends of the link must be dynamic LACP.) • **Static LACP** — Use the manually configured static LACP trunk when: – The port on the other end of the trunk link is configured for a static LACP trunk – You want to configure non-default spanning tree (STP) or IGMP parameters on an LACP trunk group. – *You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. (Refer to "VLANs and Dynamic LACP" on page 9-29.)* – You want to use a monitor port on the switch to monitor an LACP trunk. See "Trunk Group Operation Using LACP" on page 9-25.
Trunk (non-protocol)	Provides manually configured, static-only trunking to: • Most HP switches and routing switches not running the 802.3ad LACP protocol. • Windows NT and HP-UX workstations and servers Use the Trunk option when: – The device to which you want to create a trunk link is using a non-802.3ad trunking protocol – You are unsure which type of trunk to use, or the device to which you want to create a trunk link is using an unknown trunking protocol. – You want to use a monitor port on the switch to monitor traffic on a trunk. See "Trunk Group Operation Using the "Trunk" Option" on page 9-30.
FEC	Provides static trunking to forwarding devices that also support FEC (Fast EtherChannel®), such as some Cisco® switches and routers, and some HP-UX and Windows NT servers. See "Trunk Operation Using the FEC Option" on page 9-30.

Table 9-5. General Operating Rules for Port Trunks

Media: All ports on both ends of a trunk group must have the same media type and mode (speed and duplex). The switch blocks any trunked links that do not conform to this rule. (For the Series 5300XL Switches, HP recommends leaving the port Mode setting at **Auto** or, in networks using Cat 3 cabling, **Auto-10**.)

Port Configuration: The default port configuration on the Series 5300XL Switches is Auto, which enables a port to sense speed and negotiate duplex with an Auto-enabled port on another device. HP recommends that you use the Auto setting for all ports you plan to use for trunking. Otherwise, you must manually ensure that the mode setting for each port in a trunk is compatible with the other ports in the trunk.

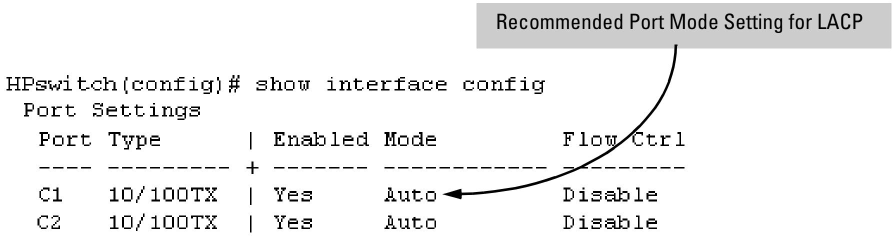

Recommended Port Mode Setting for LACP

```
HPswitch(config)# show interface config
 Port Settings
  Port Type      | Enabled Mode         Flow Ctrl
  ---- --------- + ------- ------------ ---------
  C1   10/100TX  | Yes     Auto         Disable
  C2   10/100TX  | Yes     Auto         Disable
```

All of the following operate on a per-port basis, regardless of trunk membership:

- Enable/Disable
- Flow control (Flow Ctrl)

LACP is a full-duplex protocol. See “Trunk Group Operation Using LACP” on page 9-25.

Trunk Configuration: All ports in the same trunk group must be the same trunk type (LACP, Trunk, or FEC). All LACP ports in the same trunk group must be either all static LACP or all dynamic LACP.

A trunk appears as a single port labeled **Dyn1** (for an LACP dynamic trunk) or **Trk1** (for a static trunk of any type: LACP, Trunk, or FEC) on various menu and CLI screens. For a listing of which screens show which trunk types, see “How the Switch Lists Trunk Data” on page 9-31.

For STP or VLAN operation, configuration for all ports in a trunk is done at the trunk level. (You cannot separately configure individual ports within a trunk for STP or VLAN operation.)

Traffic Distribution: All of the switch trunk protocols use the SA/DA (Source Address/Destination Address) method of distributing traffic across the trunked links. See “Outbound Traffic Distribution Across Trunked Links” on page 9-31.

Spanning Tree: Spanning Tree operates as a global setting on the switch (one instance of Spanning Tree per switch). However, you can adjust Spanning Tree parameters on a per-port basis. A static trunk of any type appears in the Spanning Tree configuration display, and you can configure Spanning Tree parameters for a static trunk in the same way that you would configure Spanning Tree parameters on a non-trunked port. (Note that the switch lists the trunk by name—such as **Trk1**—and does not list the individual ports in the trunk.) For example, if ports C1 and C2 are configured as a static trunk named **Trk1**, they are listed in the Spanning Tree display as **Trk1** and do not appear as individual ports in the Spanning Tree displays.

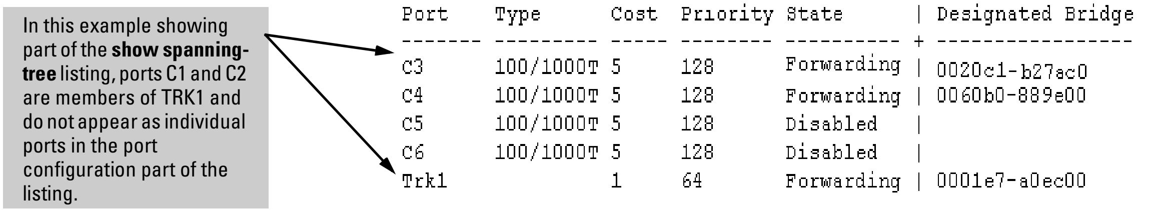

In this example showing part of the **show spanning-tree** listing, ports C1 and C2 are members of TRK1 and do not appear as individual ports in the port configuration part of the listing.

```
Port    Type      Cost  Priority State      | Designated Bridge
------- --------- ----- -------- ---------- + -----------------
C3      100/1000T 5     128      Forwarding | 0020c1-b27ac0
C4      100/1000T 5     128      Forwarding | 0060b0-889e00
C5      100/1000T 5     128      Disabled   |
C6      100/1000T 5     128      Disabled   |
Trk1              1     64       Forwarding | 0001e7-a0ec00
```

When Spanning Tree forwards on a trunk, all ports in the trunk will be forwarding. Conversely, when Spanning Tree blocks a trunk, all ports in the trunk are blocked.

Note: A dynamic LACP trunk operates only with the default Spanning Tree settings and does not appear in the Spanning Tree configuration display or **show ip igmp** listing.

If you remove a port from a static trunk, the port retains the same Spanning Tree settings that were configured for the trunk.

IP Multicast Protocol (IGMP): A static trunk of any type appears in the IGMP configuration display, and you can configure IGMP for a static trunk in the same way that you would configure IGMP on a non-trunked port. (Note that the switch lists the trunk by name—such as **Trk1**—and does not list the individual ports in the trunk.) Also, creating a new trunk automatically places the trunk in IGMP Auto status if IGMP is enabled for the default VLAN.   A dynamic LACP trunk operates only with the default IGMP settings and does not appear in the IGMP configuration display or **show ip igmp** listing.

VLANs: Creating a new trunk automatically places the trunk in the DEFAULT_VLAN, regardless of whether the ports in the trunk were in another VLAN. Similarly, removing a port from a trunk group automatically places the port in the default VLAN. You can configure a static trunk in the same way that you configure a port for membership in any VLAN.

Note: For a dynamic trunk to operate in a VLAN other than the default VLAN (DEFAULT_VLAN), GVRP must be enabled. See “Trunk Group Operation Using LACP” on page 9-25.

Port Security: Trunk groups (and their individual ports) cannot be configured for port security, and the switch excludes trunked ports from the **show port-security** listing. If you configure non-default port security settings for a port, then subsequently try to place the port in a trunk, you will see the following message and the command will not be executed:

```
< port-list > Command cannot operate over a logical port.
```

Monitor Port:

Note: A trunk cannot be a monitor port. A monitor port can monitor a static trunk but cannot monitor a dynamic LACP trunk.

Menu: Viewing and Configuring a Static Trunk Group

Important

Configure port trunking *before* you connect the trunked links to another switch, routing switch, or server. Otherwise, a broadcast storm could occur. (If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. See "Using the CLI To Configure Ports" on page 9-8.)

To View and/or Configure Static Port Trunking: This procedure uses the Port/Trunk Settings screen to configure a static port trunk group on the switch.

1. Follow the procedures in the Important note above.
2. From the Main Menu, Select:

   2. Switch Configuration . . .

   2. Port/Trunk Settings
3. Press E (for **Edit**) and then use the arrow keys to access the port trunk parameters.

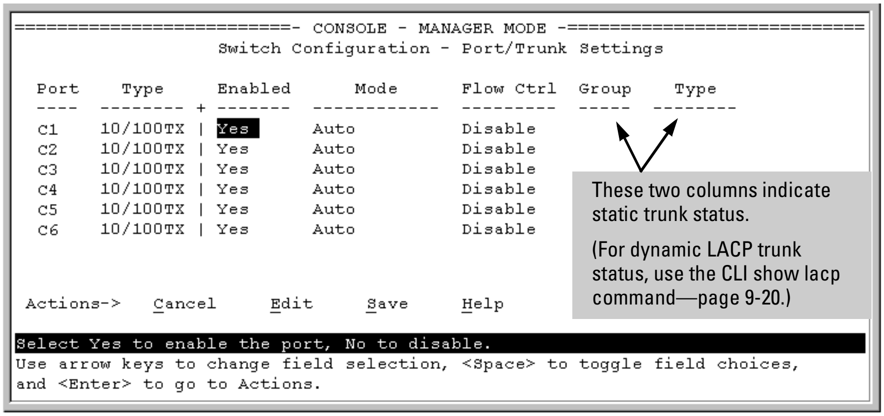

```
==========================- CONSOLE - MANAGER MODE -===========================
                   Switch Configuration - Port/Trunk Settings

  Port    Type      Enabled      Mode       Flow Ctrl  Group    Type
  ----  -------- + -------  ------------  ---------  -----  --------
  C1    10/100TX | Yes       Auto           Disable
  C2    10/100TX | Yes       Auto           Disable
  C3    10/100TX | Yes       Auto           Disable
  C4    10/100TX | Yes       Auto           Disable
  C5    10/100TX | Yes       Auto           Disable
  C6    10/100TX | Yes       Auto           Disable


 Actions->   Cancel     Edit     Save     Help

Select Yes to enable the port, No to disable.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

These two columns indicate static trunk status.

(For dynamic LACP trunk status, use the CLI show lacp command—page 9-20.)

Figure 9-6. Example of the Menu Screen for Configuring a Port Trunk Group

4. In the Group column, move the cursor to the port you want to configure.
5. Use the Space bar to choose a trunk group (**Trk1 . . . Trk36**) trunk group assignment for the selected port.

- All ports in a trunk must have the same media type and mode (such as 10/100TX set to 100FDx, or 100FX set to 100FDx). The flow control settings must also be the same for all ports in a given trunk. To verify these settings, see "Viewing Port Status and Configuring Port Parameters" on page 9-2.
- You can configure the trunk group with one, two, three, or four ports per trunk. If multiple VLANs are configured, all ports within a trunk will be assigned to the same VLAN or set of VLANs. (With the 802.1Q VLAN capability built into the switch, more than one VLAN can be assigned to a trunk. See "Port-Based Virtual LANs (Static VLANs)" on page 11-3.)

  (To return a port to a non-trunk status, keep pressing the Space bar until a blank appears in the highlighted Group value for that port.)

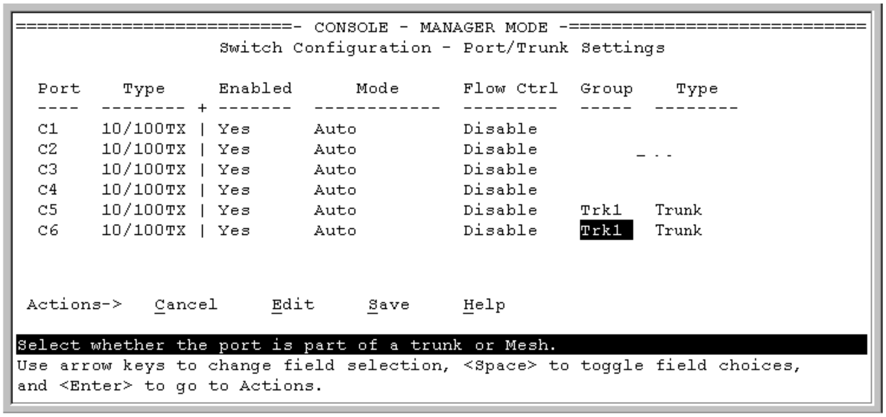

```
=========================- CONSOLE - MANAGER MODE -============================
                   Switch Configuration - Port/Trunk Settings

  Port    Type      Enabled      Mode       Flow Ctrl  Group    Type
  ----  -------- + -------  ------------  ---------  -----  --------
  C1    10/100TX | Yes      Auto          Disable
  C2    10/100TX | Yes      Auto          Disable
  C3    10/100TX | Yes      Auto          Disable
  C4    10/100TX | Yes      Auto          Disable
  C5    10/100TX | Yes      Auto          Disable    Trk1   Trunk
  C6    10/100TX | Yes      Auto          Disable    Trk1   Trunk



 Actions->   Cancel     Edit     Save     Help

Select whether the port is part of a trunk or Mesh.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Figure 9-7. Example of the Configuration for a Two-Port Trunk Group

6. Move the cursor to the Type column for the selected port and use the Space bar to select the trunk type:
   – LACP
   – Trunk (the default type if you do not specify a type)
   – FEC (Fast EtherChannel® trunk)

   All ports in the same trunk group on the same switch must have the same Type (**LACP**, **Trunk**, or **FEC**).

7. When you are finished assigning ports to the trunk group, press [Enter], then [S] (for **Save**) and return to the Main Menu. (It is not necessary to reboot the switch.)

During the Save process, traffic on the ports configured for trunking will be delayed for several seconds. If the Spanning Tree Protocol is enabled, the delay may be up to 30 seconds.

8. Connect the trunked ports on the switch to the corresponding ports on the opposite device. If you previously disabled any of the trunked ports on the switch, enable them now. (See "Viewing Port Status and Configuring Port Parameters" on page 9-2.)

Check the Event Log ("Using the Event Log To Identify Problem Sources" on page C-22) to verify that the trunked ports are operating properly.

CLI: Viewing and Configuring Port Trunk Groups

Trunk Status and Configuration Commands

show trunks	below
show lacp	page 9-20
trunk	page 9-22
interface lacp	page 9-23

Using the CLI To View Port Trunks

You can list the trunk type and group for all ports on the switch or for selected ports. You can also list LACP-only status information for LACP-configured ports.

Listing Static Trunk Type and Group for All Ports or Selected Ports.

Syntax: show trunks [<*port-list*>]

Omitting the **<** ***port-list*** **>** parameter results in a static trunk data listing for all LAN ports in the switch. For example, in a switch where ports A4 and A5 belong to Trunk 1 and ports A7 and A8 belong to Trunk 2, you have the options shown in figures 9-9-8 and 9-9-9 for displaying port data for ports belonging to static trunks.

Using a port list specifies, for switch ports in a static trunk group, only the ports you want to view. In this case, the command specifies ports A5 through A7. However, because port A6 is not in a static trunk group, it does not appear in the resulting listing:

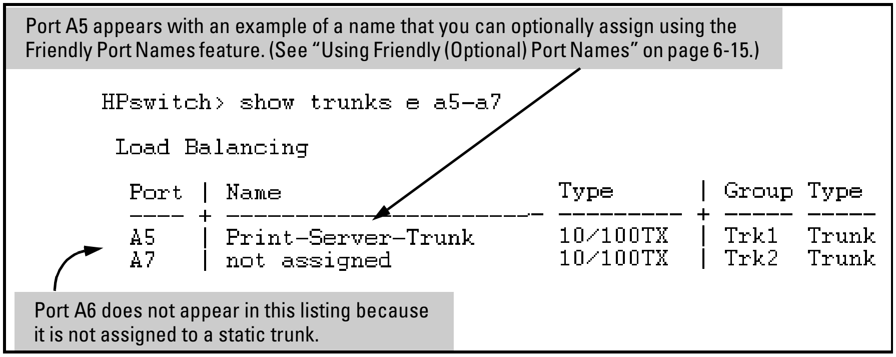

Port A5 appears with an example of a name that you can optionally assign using the Friendly Port Names feature. (See "Using Friendly (Optional) Port Names" on page 6-15.)

```
HPswitch> show trunks e a5-a7

 Load Balancing

  Port | Name                    Type      | Group Type
  ---- + ----------------------- --------- + ----- -----
  A5   | Print-Server-Trunk      10/100TX  | Trk1  Trunk
  A7   | not assigned            10/100TX  | Trk2  Trunk
```

Port A6 does not appear in this listing because it is not assigned to a static trunk.

Figure 9-8. Example Listing Specific Ports Belonging to Static Trunks

The **show trunks [e] <** ***port-list*** **>** command in the above example includes a port list, and thus shows trunk group information only for specific ports that have membership in a static trunk. In figure 9-9-9, the command does not include a port list, so the switch lists all ports having static trunk membership.

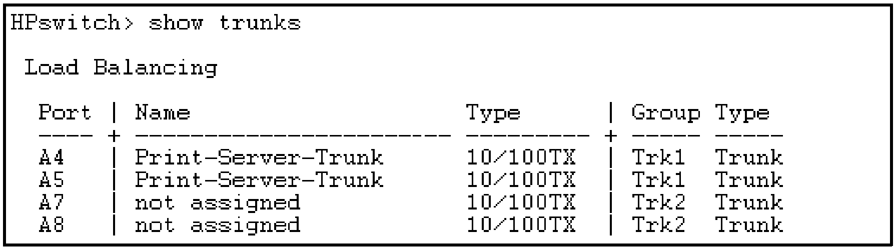

```
HPswitch> show trunks

 Load Balancing

  Port | Name                    Type      | Group Type
  ---- + ----------------------- --------- + ----- -----
  A4   | Print-Server-Trunk      10/100TX  | Trk1  Trunk
  A5   | Print-Server-Trunk      10/100TX  | Trk1  Trunk
  A7   | not assigned            10/100TX  | Trk2  Trunk
  A8   | not assigned            10/100TX  | Trk2  Trunk
```

Figure 9-9. Example of a Show Trunk Listing Without Specifying Ports

Listing Static LACP and Dynamic LACP Trunk Data. This command lists data for only the LACP-configured ports.

Syntax: show lacp

In the following example, ports A1 and A2 have been previously configured for a static LACP trunk. (For more on “Active”, see table 11-7 on page 9-27.)

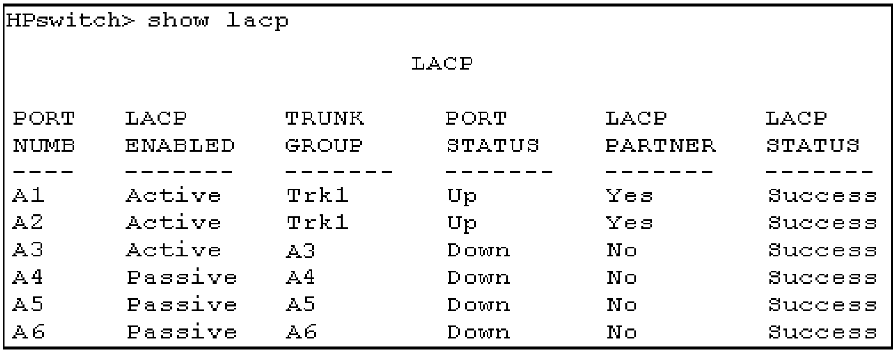

```
HPswitch> show lacp

                              LACP

 PORT    LACP       TRUNK      PORT       LACP       LACP
 NUMB    ENABLED    GROUP      STATUS     PARTNER    STATUS
 ----    -------    -------    -------    -------    -------
 A1      Active     Trk1       Up         Yes        Success
 A2      Active     Trk1       Up         Yes        Success
 A3      Active     A3         Down       No         Success
 A4      Passive    A4         Down       No         Success
 A5      Passive    A5         Down       No         Success
 A6      Passive    A6         Down       No         Success
```

Figure 9-10. Example of a Show LACP Listing

Dynamic LACP Standby Links. Dynamic LACP trunking enables you to configure standby links for a trunk by including more than four ports in a dynamic LACP trunk configuration. When four ports (trunk links) are up, the remaining link(s) will be held in standby status. If a trunked link that is “Up” fails, it will be replaced by a standby link, which maintains your intended bandwidth for the trunk. (See also the “Standby” entry under “Port Status” in "Table 9-7. LACP Port Status Data" on page 9-27.) In the next example, ports A1 through A5 have been configured for the same LACP trunk. Notice that one of the links shows Standby status, while the remaining four links are “Up”.

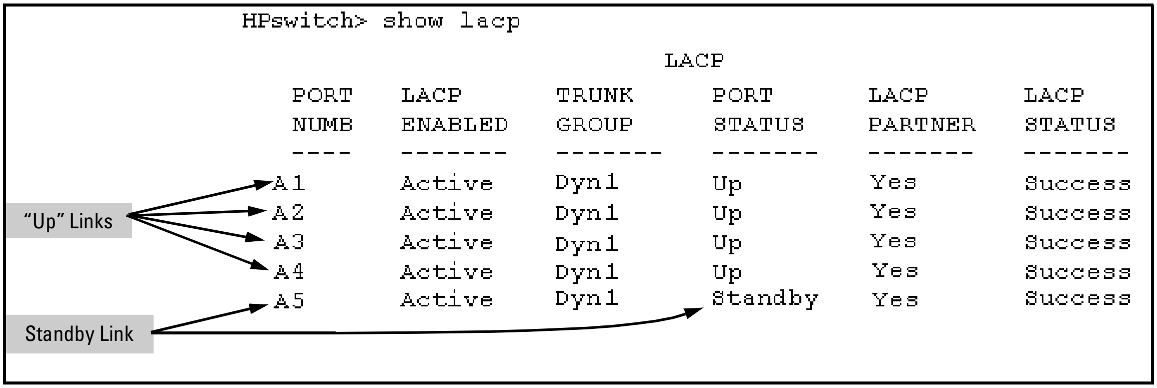

```
HPswitch> show lacp
                                   LACP
  PORT    LACP      TRUNK     PORT      LACP      LACP
  NUMB    ENABLED   GROUP     STATUS    PARTNER   STATUS
  ----    -------   -------   -------   -------   -------
 A1       Active    Dyn1      Up        Yes       Success
 A2       Active    Dyn1      Up        Yes       Success
 A3       Active    Dyn1      Up        Yes       Success
 A4       Active    Dyn1      Up        Yes       Success
 A5       Active    Dyn1      Standby   Yes       Success
```

Figure 9-11. Example of a Dynamic LACP Trunk with One Standby Link

Using the CLI To Configure a Static or Dynamic Trunk Group

Important

Configure port trunking *before* you connect the trunked links between switches. Otherwise, a broadcast storm could occur. (If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. See "Using the CLI To Configure Ports" on page 9-8.)

On the Series 5300XL Switches you can configure up to 36 port trunk groups having up to four links each (with additional standby links if you're using LACP). You can configure trunk group types as follows:

Trunk Type	**Trunk Group Membership**	
	Trk*X* (Static)	Dyn*X* (Dynamic)
LACP	Yes	Yes
Trunk	Yes	No
FEC	Yes	No

The following examples show how to create different types of trunk groups.

Configuring a Static Trunk, Static FEC, or Static LACP Trunk Group.

Syntax: trunk < trk1 | trk2 | trk3 | trk4 | trk5 | trk6 > < trunk | fec | lacp > <*port-list*>

This example uses ports C4 - C6 to create a non-protocol static trunk group with the group name of **Trk2**.

```
HPswitch(config)# trunk trk2 trunk c4-c6
```

Removing Ports from a Static Trunk Group. This command removes one or more ports from an existing Trk***x*** trunk group.

Caution

Removing a port from a trunk can result in a loop and cause a broadcast storm. When you remove a port from a trunk where STP is not in use, HP recommends that you first disable the port or disconnect the link on that port.

Syntax: no trunk < *port-list* >

This example removes ports C4 and C5 from an existing trunk group.

```
HPswitch(config)# no trunk c4-c5
```

Enabling a Dynamic LACP Trunk Group. In the default port configuration, all ports on the switch are set to LACP **Passive**. However, to enable the switch to automatically form a trunk group that is dynamic on both ends of the link, the ports on one end of a set of links must be LACP **Active**. The ports on the other end can be either LACP **Active** or LACP **Passive**. This command enables the switch to automatically establish a (dynamic) LACP trunk group when the device on the other end of the link is configured for LACP **Passive**.

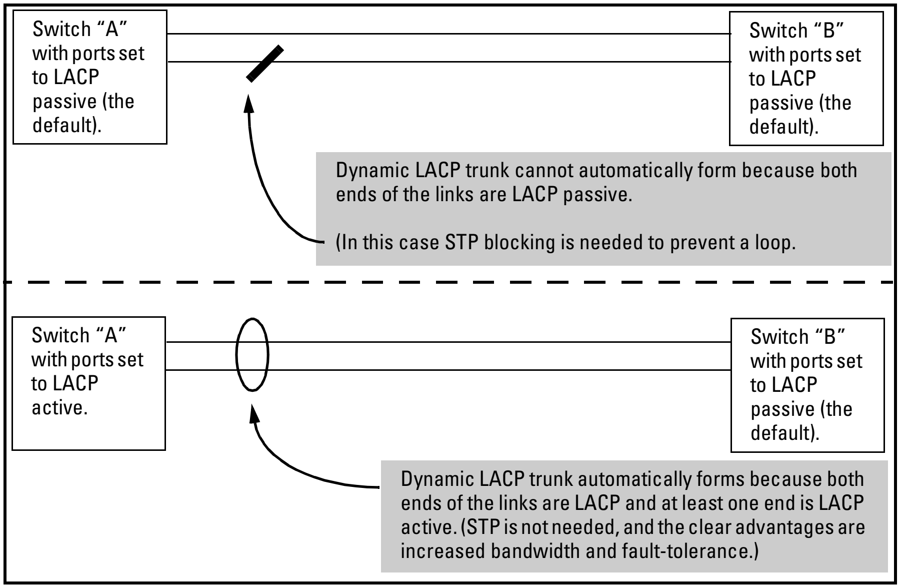
Switch “A” with ports set to LACP passive (the default).
Switch “B” with ports set to LACP passive (the default).
Dynamic LACP trunk cannot automatically form because both ends of the links are LACP passive.
(In this case STP blocking is needed to prevent a loop.
Switch “A” with ports set to LACP active.
Switch “B” with ports set to LACP passive (the default).
Dynamic LACP trunk automatically forms because both ends of the links are LACP and at least one end is LACP active. (STP is not needed, and the clear advantages are increased bandwidth and fault-tolerance.)

Figure 9-12. Example of Criteria for Automatically Forming a Dynamic LACP Trunk

Syntax: interface < *port-list* > lacp active

This example uses ports C4 and C5 to enable a dynamic LACP trunk group.

```
HPswitch(config)# interface c4-c5 lacp active
```

Removing Ports from an Dynamic LACP Trunk Group. To remove a port from dynamic LACP trunk operation, you must turn off LACP on the port. (On a port in an operating, dynamic LACP trunk, you cannot change between LACP **Active** and LACP **passive** without first removing LACP operation from the port.)

Caution

Unless STP is running on your network, removing a port from a trunk can result in a loop. To help prevent a broadcast storm when you remove a port from a trunk where STP is not in use, HP recommends that you first disable the port or disconnect the link on that port.

Syntax: no interface *<port-list>* lacp

In this example, port C6 belongs to an operating, dynamic LACP trunk. To remove port C6 from the dynamic trunk and return it to passive LACP, you would do the following:

```
HPswitch>(config)# no interface c6 lacp
HPswitch>(config)# interface c6 lacp passive
```

Note that in the above example, if the port on the other end of the link is configured for active LACP or static LACP, the trunked link will be re-established almost immediately.

Web: Viewing Existing Port Trunk Groups

While the web browser interface does not enable you to configure a port trunk group, it does provide a view of an existing trunk group.

To view any port trunk groups:

Click on the **Status** tab.

Click on [Port Status].

Trunk Group Operation Using LACP

The switch can automatically configure a dynamic LACP trunk group or you can manually configure a static LACP trunk group.

Note

LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, etc.) and the same speed, and enforces speed and duplex conformance across a trunk group.

LACP trunk status commands include:

Trunk Display Method	Static LACP Trunk	Dynamic LACP Trunk
CLI **show lacp** command	Included in listing.	Included in listing.
CLI **show trunk** command	Included in listing.	Not included.
Port/Trunk Settings screen in menu interface	Included in listing.	Not included

Thus, to display a listing of dynamic LACP trunk ports, you must use the **show lacp** command.

Note

Dynamic LACP trunks operate only in the default VLAN (unless GVRP is enabled and **Forbid** is used to prevent the trunked ports from joining the default VLAN). Thus, if an LACP dynamic port forms using ports that are not in the default VLAN, the trunk will automatically move to the default VLAN unless GVRP operation is configured to prevent this from occurring. In some cases, this can create a traffic loop in your network. For more on this topic, refer to "VLANs and Dynamic LACP" on page 9-29.

In most cases, trunks configured for LACP on the Series 5300XL Switches operate as described in table 9-6 on the next page.

Table 9-6. LACP Trunk Types

LACP Port Trunk Configuration	Operation
Dynamic LACP	This option automatically establishes an 802.3ad-compliant trunk group, with **LACP** for the port Type parameter and ***DynX*** for the port Group name, where ***X*** is an automatically assigned value from 1 to 36, depending on how many dynamic and static trunks are currently on the switch. (The switch allows a maximum of 36 trunk groups in any combination of static and dynamic trunks.) Under the following conditions, the switch automatically establishes a dynamic LACP port trunk group and assigns a port Group name: • The ports on both ends of a link have compatible mode settings (speed and duplex). • The port on one end of a link must be configured for LACP Active and the port on the other end of the same link must be configured for either LACP Passive (the default) or LACP Active. For example: 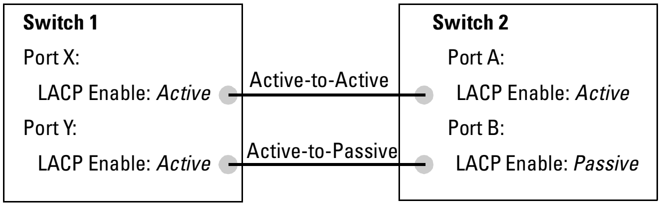 Switch 1 Port X: LACP Enable: *Active* Port Y: LACP Enable: *Active* Active-to-Active Active-to-Passive Switch 2 Port A: LACP Enable: *Active* Port B: LACP Enable: *Passive* Either of the above link configurations allow a dynamic LACP trunk link. **Standby Links:** A maximum of four operating links are allowed in the trunk, but, with dynamic LACP, you can configure one or more backup links that the switch automatically activates if a primary link fails. To configure a link as a standby for an existing four-port dynamic LACP trunk, ensure that the ports in the standby link are configured the same as either of the above examples. **Displaying Dynamic LACP Trunk Data:** To list the configuration and status for a dynamic LACP trunk, use the CLI **show lacp** command. **Note:** The dynamic trunk is automatically created by the switch, and is not listed in the static trunk listings available in the menu interface or in the CLI **show trunk** listing.
Static LACP	The trunk operates if the trunk group on the opposite device is running one of the following trunking protocols: • Active LACP • Passive LACP • Trunk • FEC This option uses **LACP** for the port Type parameter and **Trk*X*** for the port Group parameter, where ***X*** is an automatically assigned value from 1 to 36, depending on how many static trunks are currently operating on the switch. (The switch allows a maximum of six trunk groups in any combination of static and dynamic trunks.) Displaying Static LACP Trunk Data: To list the configuration and status for a static LACP trunk, use the CLI **show lacp** command. To list a static LACP trunk with its assigned ports, use the CLI **show trunk** command or display the menu interface Port/Trunk Settings screen. Static LACP does not allow standby ports.

Default Port Operation

In the default configuration, all ports are configured for passive LACP. However, if LACP is not configured, the port will not try to detect a trunk configuration and will operate as a standard, untrunked port. The following table describes the elements of per-port LACP operation. To display this data for a particular switch, execute the following command in the CLI:

```
HPswitch> show lacp
```

Table 9-7. LACP Port Status Data

Status Name	Meaning
Port Numb	Shows the physical port number for each port configured for LACP operation (C1, C2, C3 . . .). Unlisted port numbers indicate that the missing ports are assigned to a static Trunk group, an FEC trunk group, or are not configured for any trunking.
LACP Enabled	**Active:** The port automatically sends LACP protocol packets. **Passive:** The port does not automatically send LACP protocol packets, and responds only if it receives LACP protocol packets from the opposite device. A link having either two active LACP ports or one active port and one passive port can perform dynamic LACP trunking. A link having two passive LACP ports will not perform LACP trunking because both ports are waiting for an LACP protocol packet from the opposite device. **Note:** In the default switch configuration, all ports are configured for passive LACP operation.
Trunk Group	**Trk*X***: This port has been manually configured into a static LACP trunk. **Trunk Group Same as Port Number:** The port is configured for LACP, but is not a member of a port trunk.
Port Status	**Up:** The port has an active LACP link and is not blocked or in Standby mode. **Down:** The port is enabled, but an LACP link is not established. This can indicate, for example, a port that is not connected to the network or a speed mismatch between a pair of linked ports. **Disabled:** The port cannot carry traffic. **Blocked:** LACP, STP, or FEC has blocked the port. (The port is not in LACP Standby mode.) This may be due to a trunk negotiation (very brief) or a configuration error such as differing port speeds on the same link or attempting to connect the Series 5300XL Switches to more than 36 trunks. **Standby:** The port is configured for dynamic LACP trunking to another device, but the maximum number of ports for the Dynamic trunk to that device has already been reached on either the Series 5300XL Switches or the other device. This port will remain in reserve, or "standby" unless LACP detects that another, active link in the trunk has become disabled, blocked, or down. In this case, LACP automatically assigns a Standby port, if available, to replace the failed port.
LACP Partner	**Yes:** LACP is enabled on both ends of the link. **No:** LACP is enabled on the Series 5300XL Switches, but either LACP is not enabled or the link has not been detected on the opposite device.
LACP Status	**Success:** LACP is enabled on the port, detects and synchronizes with a device on the other end of the link, and can move traffic across the link. **Failure:** LACP is enabled on a port and detects a device on the other end of the link, but is not able to synchronize with this device, and therefore not able to send LACP packets across the link. This can be caused, for example, by an intervening device on the link (such as a hub), a bad hardware connection, or if the LACP operation on the opposite device does not comply with the IEEE 802.3ad standard.

LACP Notes and Restrictions

802.1x (Port-Based Access Control) Configured on a Port. To maintain security, LACP is not allowed on ports configured for 802.1x authenticator operation. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables 802.1x on that port.

```
HPswitch(config)# aaa port-access authenticator e b1
LACP has been disabled on 802.1x port(s).
HPswitch(config)#
```

The switch will not allow you to configure LACP on a port on which port access (802.1x) is enabled. For example:

```
HPswitch(config)# int e b1 lacp passive
Error configuring port < port-number >: LACP and 802.1x
cannot be run together.
HPswitch(config)#
```

To restore LACP to the port, you must first remove the port's 802.1x configuration and then re-enable LACP active or passive on the port.

Port Security Configured on a Port. To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables port security on that port. For example:

```
HPswitch(config)# port-security e a17 learn-mode static
address-limit 2
LACP has been disabled on secured port(s).
HPswitch(config)#
```

The switch will not allow you to configure LACP on a port on which port security is enabled. For example:

```
HPswitch(config)# int e a17 lacp passive
Error configuring port A17: LACP and port security cannot
be run together.
HPswitch(config)#
```

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.

Changing Trunking Methods. To convert a trunk from static to dynamic, you must first eliminate the static trunk.

Static LACP Trunks. Where a port is configured for LACP (Active or Passive), but does not belong to an existing trunk group, you can add that port to a static trunk. Doing so disables dynamic LACP on that port, which means you must manually configure both ends of the trunk.

Dynamic LACP Trunks. You can configure a port for LACP-active or LACP-passive, but on a dynamic LACP trunk you cannot configure the other options that you can on static trunks. If you want to manually configure a trunk, use the **trunk** command. (Refer to "Using the CLI To Configure a Static or Dynamic Trunk Group" on page 9-21.)

VLANs and Dynamic LACP. A dynamic LACP trunk operates only in the default VLAN (unless you have enabled GVRP on the switch and use **Forbid** to prevent the ports from joining the default VLAN).

- If you want to use LACP for a trunk on a non-default VLAN and GVRP is disabled, configure the trunk as a static trunk.
- If there are ports that you do not want on the default VLAN, ensure that they cannot become dynamic LACP trunk members. Otherwise a traffic loop can unexpectedly occur. For example:

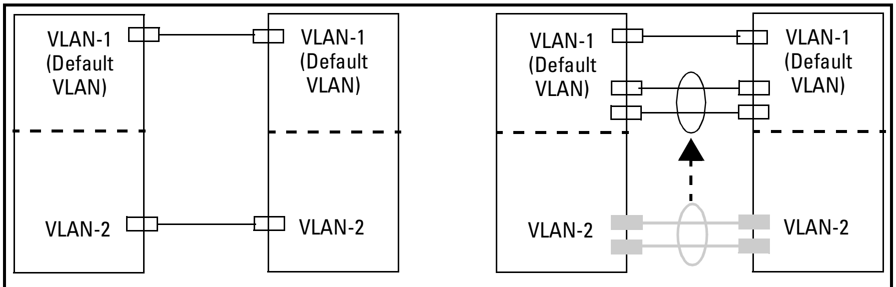

If the ports in VLAN 2 are configured to allow a dynamic trunk (and GVRP is disabled), adding a second link in VLAN 2 automatically forms a dynamic LACP trunk and moves the trunk to VLAN-1 (the default VLAN), which creates a traffic loop in VLAN 1 between the two switches and eliminates the link in VLAN 2 between the two switches.

Figure 9-13. A Dynamic LACP Trunk Forming in a VLAN Can Cause a Traffic Loop

Easy control methods include either disabling LACP on the selected ports or configuring them to operate in static LACP trunks.

Spanning Tree and IGMP. If Spanning Tree and/or IGMP is enabled in the switch, a dynamic LACP trunk operates only with the default settings for these features and does not appear in the port listings for these features.

Half-Duplex and/or Different Port Speeds Not Allowed in LACP Trunks. The ports on both sides of an LACP trunk must be configured for the same speed and for full-duplex (FDx). The 802.3ad LACP standard specifies a full-duplex (FDx) requirement for LACP trunking.

A port configured as LACP passive and not assigned to a port trunk can be configured to half-duplex (HDx). However, in any of the following cases, a port cannot be reconfigured to an HDx setting:

- If a port is set to LACP Active, you cannot configure it to HDx.
- If a port is already a member of a static or dynamic LACP trunk, you cannot configure it to HDx.
- If a port is already set to HDx, the switch does not allow you to configure it for a static or dynamic LACP trunk.

Dynamic/Static LACP Interoperation: A port configured for dynamic LACP can properly interoperate with a port configured for static (Trk*X*) LACP, but any ports configured as standby LACP links will be ignored.

Trunk Group Operation Using the "Trunk" Option

This method creates a trunk group that operates independently of specific trunking protocols and does not use a protocol exchange with the device on the other end of the trunk. With this choice, the switch simply uses the SA/DA method of distributing outbound traffic across the trunked ports without regard for how that traffic is handled by the device at the other end of the trunked links. Similarly, the switch handles incoming traffic from the trunked links as if it were from a trunked source.

Use the Trunk option when you are trying to establish a trunk group between a Series 5300XL switch and another device, but the other device's trunking operation fails to interoperate properly with LACP or FEC trunking configured on the Series 5300XL Switches.

Trunk Operation Using the "FEC" Option

This is the most flexible method for distributing traffic over trunked links when connecting to devices that use the FEC (Fast EtherChannel®) technology. FEC trunks offer the following benefits:

- Provide trunked connectivity to a FEC-compliant server, switch, or router.
- Enable quick convergence to remaining links when a failure is detected on a trunked port link.

- Depending on the capabilities of the device on the other end of the trunk, negotiate the forwarding mechanism on the trunk to the non-protocol option.
- When auto-negotiated to the SA/DA forwarding mechanism, provide higher performance on the trunk for broadcast, multicast, and flooded traffic through distribution in the same manner as non-protocol trunking.
- Support FEC automatic trunk configuration mode on other devices. That is, when connecting FEC trunks to FEC-capable servers, switches, or routers having FEC automatic trunk configuration mode enabled, the FEC trunks allow these other devices to automatically form trunk groups.

How the Switch Lists Trunk Data

Static Trunk Group: Appears in the menu interface and the output from the CLI **show trunk** and **show interfaces** commands.

Dynamic LACP Trunk Group: Appears in the output from the CLI **show lacp** command.

Interface Option	Dynamic LACP Trunk Group	Static LACP Trunk Group	Static Non-Protocol or FEC Trunk Group
Menu Interface	No	Yes	Yes
CLI **show trunk**	No	Yes	Yes
CLI **show interfaces**	No	Yes	Yes
CLI **show lacp**	Yes	Yes	No
CLI **show spanning-tree**	No	Yes	Yes
CLI **show igmp**	No	Yes	Yes
CLI **show config**	No	Yes	Yes

Outbound Traffic Distribution Across Trunked Links

All three trunk group options (LACP, Trunk, and FEC) use source-destination address pairs (SA/DA) for distributing outbound traffic over trunked links.

SA/DA (source address/destination address) causes the switch to distribute outbound traffic to the links within the trunk group on the basis of source/ destination address pairs. That is, the switch sends traffic from the same source address to the same destination address through the same trunked link, and sends traffic from the same source address to a different destination address through a different link, depending on the rotation of path assignments among the links in the trunk. Likewise, the switch distributes traffic for the same destination address but from different source addresses through different links. Because the amount of traffic coming from or going to various nodes in a network can vary widely, it is possible for one link in a trunk group to be fully utilized while others in the same trunk have unused bandwidth capacity even though the address assignments are evenly distributed across the links in a trunk. In actual networking environments, this is rarely a problem. However, if it becomes a problem, you can use the HP TopTools for Hubs & Switches network management software available from Hewlett-Packard to quickly and easily identify the sources of heavy traffic (top talkers) and make adjustments to improve performance.

Broadcasts, multicasts, and floods from different source addresses are distributed evenly across the links. As links are added or deleted, the switch redistributes traffic across the trunk group. For example, in figure 9-9-14 showing a three-port trunk, traffic could be assigned as shown in table 9-8.

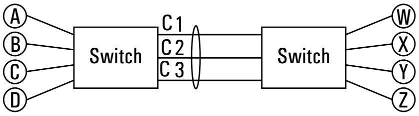

Figure 9-14. Example of Port-Trunked Network

Table 9-8. Example of Link Assignments in a Trunk Group (SA/DA Distribution)

Source:	Destination:	Link:
Node A	Node W	1
Node B	Node X	2
Node C	Node Y	3
Node D	Node Z	1
Node A	Node Y	2
Node B	Node W	3

Traffic/Security Filters

Feature	Default	Menu	CLI	Web
configure source-port filters	none	n/a	page 9-35	n/a
configure protocol filters	none	n/a	page 9-35	n/a
configure multicast filters	none	n/a	page 9-35	n/a
display filter data	n/a	n/a	page 9-36	n/a

From the CLI you can enhance in-band security and bandwidth usage by configuring static filters on the switch to either forward (the default) or drop inbound traffic meeting the filter criteria, as described in table 9-9.

Table 9-9. Filter Types and Criteria

Static Filter Type	Selection Criteria
Multicast	Inbound traffic having a specified multicast MAC address will be forwarded to outbound ports (the default) or dropped on a per-port (destination) basis.
Protocol	Inbound traffic having the selected frame (protocol) type will be forwarded or dropped on a per-port (destination) basis.
Source-Port	Inbound traffic from a designated source-port will be forwarded or dropped on a per-port (destination) basis.

Filter Limits

The switch accepts up to 101 static filters. These limitations also apply:

- Multicast filters: up to 16
- Protocol filters: up to 7
- Source-port filters: Up to 78

For configuration information, turn to the next page. For more information on filter types and operation, refer to “Filter Types and Operation” on page 9-38.

Steps for Configuring Traffic/Security Filters

Use this procedure to specify the type of filters to use on the switch and whether to forward or drop filtered packets for each filter you specify.

1. Select the static filter type(s) (source-port, protocol, and/or multicast).
2. For inbound traffic matching the filter type, determine the filter action you want for each outbound (destination) port on the switch (forward or drop). The default action for a new filter is to forward traffic of the specified type to all outbound ports.
3. Configure the filter.
4. Check the filter listing to verify that you have configured correct action for the desired outbound ports.

Configuring a Traffic Filter. The **filter** command specifies the filter type and action, and the destination (outbound) ports on which to apply the action.

Syntax: filter

[source-port [ethernet] < *port-number* >]

Specifies one inbound port. Traffic received on this port from other devices will be filtered. ***Note:*** *If multiple VLANs are configured, the source-port and the destination port(s) must be in the same VLAN. (Default: Forward on all ports.)*

[< forward | drop > [e] < *port-list* >]

Specifies whether the designated destination port(s) should forward or drop the filtered traffic.

[multicast < *MAC- address* >]

Specifies a multicast address. Inbound traffic received (on any port) with this multicast address will be filtered. (Default: Forward on all ports.)

[< forward | drop > [e] < *port-list* >]

Specifies whether the designated destination port(s) should forward or drop the filtered traffic.

[protocol < ip | ipx | arp | dec-lat | appletalk | sna | netbeui >]

Specifies a protocol type. Traffic received (on any port) with this protocol type will be filtered. (Default: Forward on all ports.)

[< forward | drop > [e] < *port-list* >]

Specifies whether the designated destination port(s) should forward or drop the filtered traffic.

no filter source-port [e] < *port-number* >

Deletes the source-port filter for < port-number > and returns the destination ports for that filter to the **Forward** *action.*

no filter multicast < *MAC-address* >

Deletes the multicast filter for the < MAC-address > multicast address and returns the destination ports for that filter to the **Forward** *action.*

no filter protocol < ip | ipx | arp | dec-lat | appletalk | sna | netbeui >

Deletes the protocol filter for the specified protocol and returns the destination ports for that filter to the **Forward** *action.*

For example, suppose you wanted to configure these six filters on the switch:

Filter Type	Filter Value	Action	Destination Ports
Source-Port	Inbound ports: A1, A2*	Drop	D1-D4
Multicast	010000-123456	Drop	C1-C24, D5-D10
Multicast	010000-224466	Drop	B1-B4
Protocol	Appletalk	Drop	C12-C18, D1
Protocol	ARP	Drop	D17, D21-D24

*Because the switch allows one inbound port in a source-port filter, the requirement to filter ports A1 and A2 means you will configure two separate source-port filters.

The following commands configure the filters listed above:

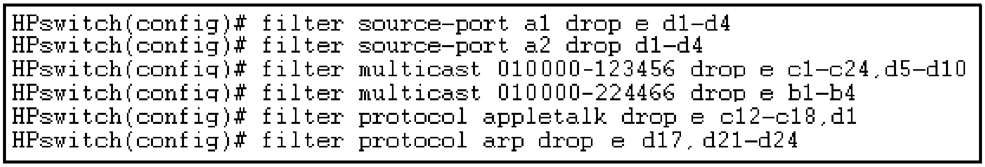

```
HPswitch(config)# filter source-port a1 drop e d1-d4
HPswitch(config)# filter source-port a2 drop d1-d4
HPswitch(config)# filter multicast 010000-123456 drop e c1-c24,d5-d10
HPswitch(config)# filter multicast 010000-224466 drop e b1-b4
HPswitch(config)# filter protocol appletalk drop e c12-c18,d1
HPswitch(config)# filter protocol arp drop e d17, d21-d24
```

Figure 9-15. Configuring Various Traffic/Security Filters

Displaying Traffic/Security Filters. This command displays a listing of all filters by index number and also enables you to use the index number to display the details of individual filters.

Syntax: show filter

Show a table listing the filters configured in the switch, with corresponding filter index numbers.

[*index*]

Lists the filter type and other relevant data for the filter corresponding to the index number. Also lists, for each outbound destination port in the switch, the port number, port type, and filter action (forward or drop). The switch assigns the lowest available index number to a new filter. If you delete a filter, the index number for that filter becomes available for the next new filter you create.

For example, to display the filters created in figure 9-15 and then list the details of the multicast filter for multicast address **010000-224466**:

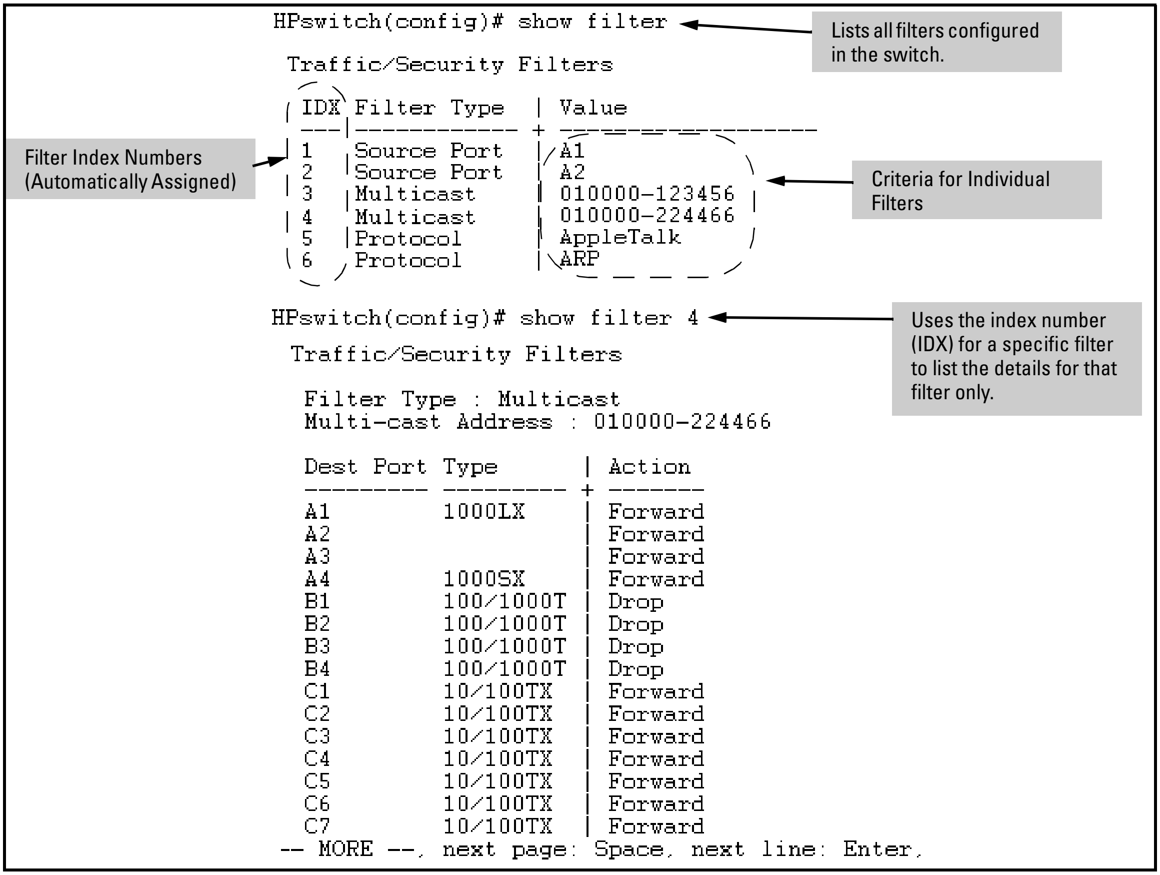

```
HPswitch(config)# show filter

 Traffic/Security Filters

  IDX Filter Type   | Value
  ---|------------- + ------------------
  1   Source Port   | A1
  2   Source Port   | A2
  3   Multicast     | 010000-123456
  4   Multicast     | 010000-224466
  5   Protocol      | AppleTalk
  6   Protocol      | ARP

HPswitch(config)# show filter 4
 Traffic/Security Filters

  Filter Type : Multicast
  Multi-cast Address : 010000-224466

  Dest Port Type       | Action
  --------- ---------- + -------
  A1        1000LX     | Forward
  A2                   | Forward
  A3                   | Forward
  A4        1000SX     | Forward
  B1        100/1000T  | Drop
  B2        100/1000T  | Drop
  B3        100/1000T  | Drop
  B4        100/1000T  | Drop
  C1        10/100TX   | Forward
  C2        10/100TX   | Forward
  C3        10/100TX   | Forward
  C4        10/100TX   | Forward
  C5        10/100TX   | Forward
  C6        10/100TX   | Forward
  C7        10/100TX   | Forward
 -- MORE --, next page: Space, next line: Enter,
```

Lists all filters configured in the switch.

Filter Index Numbers (Automatically Assigned)

Criteria for Individual Filters

Uses the index number (IDX) for a specific filter to list the details for that filter only.

Figure 9-16. Example of Displaying Filter Data

Filter Types and Operation

Multicast Filters

This filter type enables the switch to forward or drop multicast traffic to a specific set of destination ports. This helps to preserve bandwidth by reducing multicast traffic on ports where it is unnecessary, and to isolate multicast traffic to enhance security.

You can configure up to 16 static multicast filters (defined by the **filter** command). However, IGMP-controlled filters will override any static multicast filters having the same multicast address as specified by IGMP. The total of static multicast filters and IGMP multicast filters together can range from 389 to 420, depending on the current **max-vlans** setting in the switch.

Table 9-10. Maximum Allowed Number of Multicast Filters

Max-VLANs Setting	Maximum # of Multicast Filters (Static and IGMP Combined)
1 (the minimum)	420
8 (the default)	413
32 or higher	389

Note:

IP Multicast Filters. Multicast filters are configured using the Ethernet format for the multicast address. IP multicast addresses occur in the range from 224.0.0.0 through 239.255.255.255 (which corresponds to the Ethernet multicast address range of 01005e-000000 through 01005e-7fffff). Any static Traffic/Security filters configured with a **multicast** filter type and a multicast address in this range will continue to be in effect unless IGMP learns of a multicast group destination in this range. In this case, IGMP takes over the filtering function for the multicast destination address(es) for as long as the IGMP group is active. If the IGMP group subsequently deactivates, the static filter resumes control over traffic to the multicast address.

Caution

If Spanning Tree is enabled, then the Spanning Tree multicast MAC address should not be filtered. (STP will not operate properly if the STP multicast MAC address is filtered.)

Protocol Filters

This filter type enables the switch to forward or drop, on the basis of protocol type, traffic to a specific set of destination ports on the switch. Filtered protocol types include:

- AppleTalk
- ARP
- DEC LAT
- IP
- IPX
- NetBEUI
- SNA

Only one filter for a particular protocol type can be configured at any one time. For example, a separate protocol filter can be configured for each of the protocol types listed above, but only one of those can be an IP filter. Also, the destination ports for a protocol filter can be on different VLANs.

You can configure up to seven protocol filters.

Source-Port Filters

This filter type enables the switch to forward or drop traffic from *all* end nodes on the indicated source-port to specific destination ports. If VLANs are configured, the destination port must be in the same VLAN as the source-port. Only one source-port filter can be configured for each of the ports in the switch.

You can configure all destination ports in the switch on a single source-port filter.

Note

If more than one VLAN is configured, then the set of destination ports (Dest Port parameter) can consist of only the destination ports that are in the same VLAN as the source-port.

10

Configuring for Network Management Applications

Contents

Using SNMP Tools To Manage the Switch

Overview

You can manage the switch via SNMP from a network management station, including HP TopTools for Hubs & Switches — an OpenView-based network management application that runs on your Windows NT- or Windows 2000-based PC. HP TopTools for Hubs & Switches provides control of your switch through its web browser interface. In addition, it uses the RMON agent statistical sampling software that is included in the switch to provide easy-to-use traffic monitoring and network activity analysis tools. For more on TopTools, see the "Read Me First" document shipped with your switch and also available on HP's Procurve web site at

http://www.hp.com/go/procurve

This section includes:

- An overview of SNMP management for the switch
- Configuring the switches for:
  - SNMP Communities (page 10-4)
  - Trap Receivers and Authentication Traps (page 10-8)
- Information on advanced management through RMON Support (page 10-12)

To implement SNMP management, the switch must have an IP address, configured either manually or dynamically (with DHCP or Bootp). If multiple VLANs are configured, each VLAN interface should have its own IP address. For DHCP use with multiple VLANs, see "The Primary VLAN" on page 11-6.

Note

If you use the switch's Authorized IP Managers and Management VLAN features, ensure that the SNMP management station and/or the choice of switch port used for SNMP access to the switch are compatible with the access controls enforced by these features. Otherwise, SNMP access to the switch will be blocked. For more on Authorized IP Managers, refer to the Access Security Guide on the Documentation CD-ROM shipped with your switch and also available on the HP Procurve web site. For information on the Management VLAN feature, refer to "The Secure Management VLAN" on page 11-27.

SNMP Management Features

SNMP management features on the switch include:

- SNMP version 2c over IP
- Security via configuration of SNMP communities (page 10-3)
- Event reporting via SNMP
  - Version 1 traps
  - RMON: groups 1, 2, 3, and 9
- Managing the switch with an SNMP network management tool such as HP TopTools for Hubs & Switches
- Standard MIBs, such as the Bridge MIB (RFC 1493), Ethernet MAU MIB (RFC 1515), and others.

The switch SNMP agent also uses certain variables that are included in a Hewlett-Packard proprietary MIB (Management Information Base) file. To ensure that you have the latest version in the database of your TopTools, OpenView, or other SNMP network management tool, you can copy the MIB file from the HP Procurve World Wide Web site at:

http://www.hp.com/go/hpprocurve

Click on **software**, then **MIBs**.

Configuring for SNMP Access to the Switch

SNMP access requires an IP address and subnet mask configured on the switch. (See “IP Configuration” on page 7-3.) If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address. (See “DHCP/Bootp Operation” on page 7-13.)

Once an IP address has been configured, the main steps for configuring for SNMP access to management features are:

1. Configure the appropriate SNMP communities. (Refer to “SNMP Communities” on page 10-4.)
2. Configure the appropriate trap receivers. (Refer to “Trap Receivers and Authentication Traps” on page 10-8.)

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct community name may access the switch with the View and Access levels that have been set for that community. If you want to restrict access to one or more specific nodes, you can use the switch’s IP Authorized Manager feature. (Refer to the *Access Security Guide* for your switch.)

Caution The “public” community exists by default and is used by HP’s network management applications. Deleting the “public” community disables many network management functions (such as auto-discovery, traffic monitoring, SNMP trap generation, and threshold setting). If security for network management is a concern, it is recommended that you change the write access for the “public” community to “Restricted”.

SNMP Communities

SNMP Community Features

Feature	Default	Menu	CLI	Web
show SNMP communities	n/a	page 10-4	page 10-6	—
configure identity information	none	—	page 10-7	
configure community names	public	page 10-4	page 10-7	—
MIB view for a community name (operator, manager)	manager	“	“	
write access for default community name	unrestricted	“	“	

Use SNMP communities to restrict access to the switch by SNMP management stations by adding, editing, or deleting SNMP communities. You can configure up to five SNMP communities, each with either an operator-level or a manager-level view, and either restricted or unrestricted write access.

Using SNMP requires that the switch have an IP address and subnet mask compatible with your network.

Caution Deleting or changing the community named “public” prevents network management applications (such as auto-discovery, traffic monitoring, SNMP trap generation, and threshold setting) from operating in the switch. (Changing or deleting the “public” name also generates an Event Log message.) If security for network management is a concern, it is recommended that you change the write access for the “public” community to “Restricted”.

Menu: Viewing and Configuring SNMP Communities

To View, Edit, or Add SNMP Communities:

1. From the Main Menu, Select:

   2. Switch Configuration...

   6. SNMP Community Names

Note: This screen gives an overview of the SNMP communities that are currently configured. All fields in this screen are read-only.

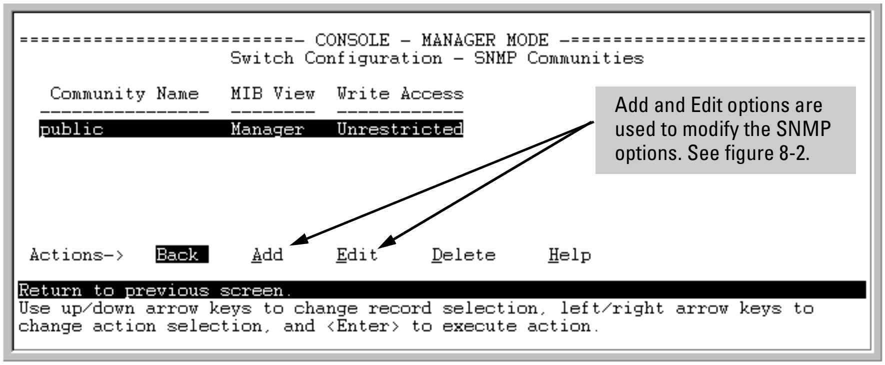

```
=========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - SNMP Communities

  Community Name    MIB View   Write Access
  ----------------  --------   ------------
  public            Manager    Unrestricted





 Actions->   Back      Add      Edit      Delete     Help

Return to previous screen.
Use up/down arrow keys to change record selection, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Add and Edit options are used to modify the SNMP options. See figure 8-2.

Figure 10-1. The SNMP Communities Screen (Default Values)

2. Press [A] (for **Add**) to display the following screen:

If you are adding a community, the fields in this screen are blank.

If you are editing an existing community, the values for the currently selected Community appear in the fields.

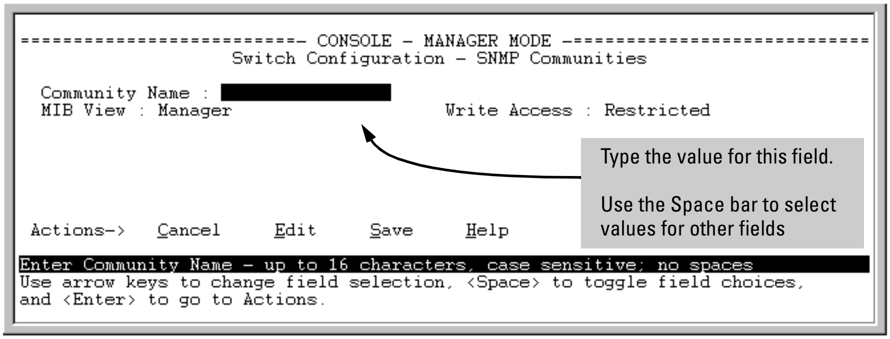

```
=========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - SNMP Communities

  Community Name :
  MIB View : Manager                    Write Access : Restricted





 Actions->   Cancel     Edit     Save     Help

Enter Community Name - up to 16 characters, case sensitive; no spaces
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Type the value for this field.

Use the Space bar to select values for other fields

Figure 10-2. The SNMP Add or Edit Screen

Need Help? If you need information on the options in each field, press [Enter] to move the cursor to the Actions line, then select the **Help** option on the Actions line. When you are finished with Help, press [E] (for Edit) to return the cursor to the parameter fields.

3. Enter the name you want in the Community Name field, and use the Space bar to select the appropriate value in each of the other fields. (Use the [Tab] key to move from one field to the next.)

4. Press [Enter], then [S] (for **Save**).

CLI: Viewing and Configuring Community Names

Community Name Commands	**Page**
show snmp-server [<*community-string*>]	10-6
[no] snmp-server	10-7
[community <*community-str*>]	10-7
[host <*community-str*> <*ip-addr*>] [<none | debug | all | not-info | critical>]	10-10
[enable traps <authentication>	10-11

Listing Community Names and Values. This command lists the data for currently configured SNMP community names (along with trap receivers and the setting for authentication traps — see “Trap Receivers and Authentication Traps” on page 10-8).

Syntax: show snmp-server [<*community-string*>]

This example lists the data for all communities in a switch; that is, both the default “public” community name and another community named “red-team”

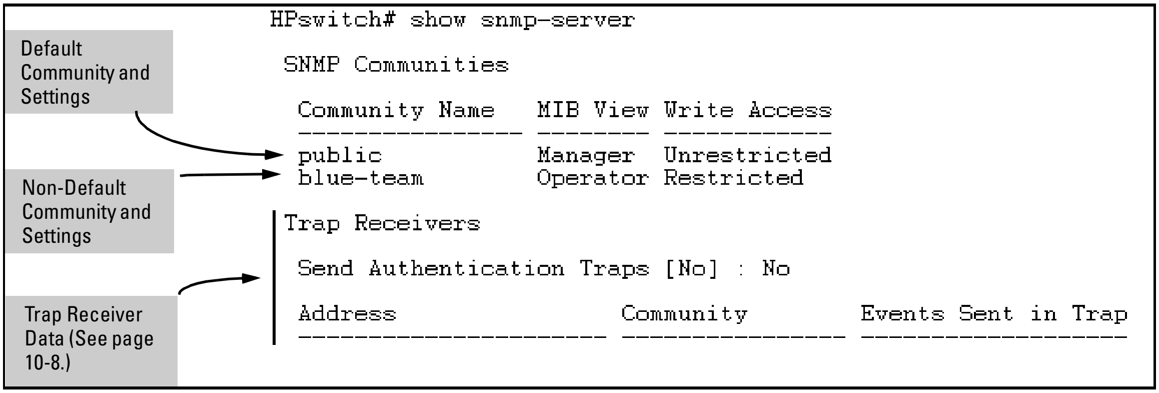

```
HPswitch# show snmp-server

 SNMP Communities

  Community Name    MIB View Write Access
  ---------------- -------- ------------
  public           Manager  Unrestricted
  blue-team        Operator Restricted

 Trap Receivers

  Send Authentication Traps [No] : No

  Address                Community        Events Sent in Trap
  ---------------------- ---------------- --------------------
```

Default Community and Settings
Non-Default Community and Settings
Trap Receiver Data (See page 10-8.)

Figure 10-3. Example of the SNMP Community Listing with Two Communities

To list the data for only one community, such as the “public” community, use the above command with the community name included. For example:

```
HPswitch# show snmp-server public
```

Configuring Community Names and Values. The **snmp-server** command enables you to add SNMP communities with either default or specific access attributes, and to delete specific communities.

Syntax: [no] snmp-server community < *community-name* >

Configures a new community name. If you do not also specify **operator** *or* **manager**, *the switch automatically assigns the community to the* **operator** *MIB view. If you do not specify* **restricted** *or* **unrestricted**, *the switch automatically assigns the community to* **restricted** *(read-only) access. The* **no** *form uses only the* ***< community-name >*** *variable and deletes the named community from the switch.*

[operator | manager]

Optionally assigns an access level. At the **operator** *level the community can access all MIB objects except the CONFIG MIB. At the* **manager** *level the community can access all MIB objects.*

[restricted | unrestricted]

Optionally assigns MIB access type. Assigning the **restricted** *type allows the community to read MIB variables, but not to set them. Assigning the* **unrestricted** *type allows the community to read and set MIB variables.*

For example, to add the following communities:

Community	Access Level	Type of Access
red-team	manager *(Access to all MIB objects.)*	unrestricted *(read/write)*
blue-team	operator *(Access to all MIB objects except the CONFIG MIB.)*	restricted *(read-only)*

```
HPswitch(config)# snmp-server community red-team
                  manager unrestricted
HPswitch(config)# snmp-server community blue-team
                  operator restricted
```

To eliminate a previously configured community named “gold-team”:

```
HPswitch(config) # no snmp-server community gold-team
```

Trap Receivers and Authentication Traps

Trap Features

Feature	Default	Menu	CLI	Web
snmp-server host (trap receiver)	public	—	page 10-10	—
snmp-server enable (authentication trap)	none	—	page 10-11	—

A *trap receiver* is a management station designated by the switch to receive SNMP traps sent from the switch. An *authentication trap* is a specialized SNMP trap sent to trap receivers when an unauthorized management station tries to access the switch.

Note

Fixed or "Well-Known" Traps: The Series 5300XL switches automatically sends fixed traps (such as "coldStart", "warmStart", "linkDown", and "linkUp") to trap receivers using the **public** community name. These traps cannot be redirected to other communities. Thus, if you change or delete the default **public** community name, these traps will be lost.

Thresholds: The switch automatically sends all messages resulting from thresholds to the network management station(s) that set the thresholds, regardless of the trap receiver configuration.

In the default configuration, there are no trap receivers configured, and the authentication trap feature is disabled. From the CLI you can configure up to ten SNMP trap receivers to receive SNMP traps from the switch. As an option, you can also configure the switch to send Event Log messages as traps. CLI: Configuring and Displaying Trap Receivers

Trap Receiver Commands	Page
show snmp-server	10-9
snmp-server host *<ip-addr> <community-name>* [none | all | non-infol critical | debug]	10-10
snmp-server enable traps authentication	10-10

Using the CLI To List Current SNMP Trap Receivers.

This command lists the currently configured trap receivers and the setting for authentication traps (along with the current SNMP community name data — see “SNMP Communities” on page 10-4).

Syntax: show snmp-server

Displays current community and trap receiver data.

In the next example, the **show snmp-server** command shows that the switch has been previously configured to send SNMP traps to management stations belonging to the “public”, “red-team”, and “blue-team” communities.

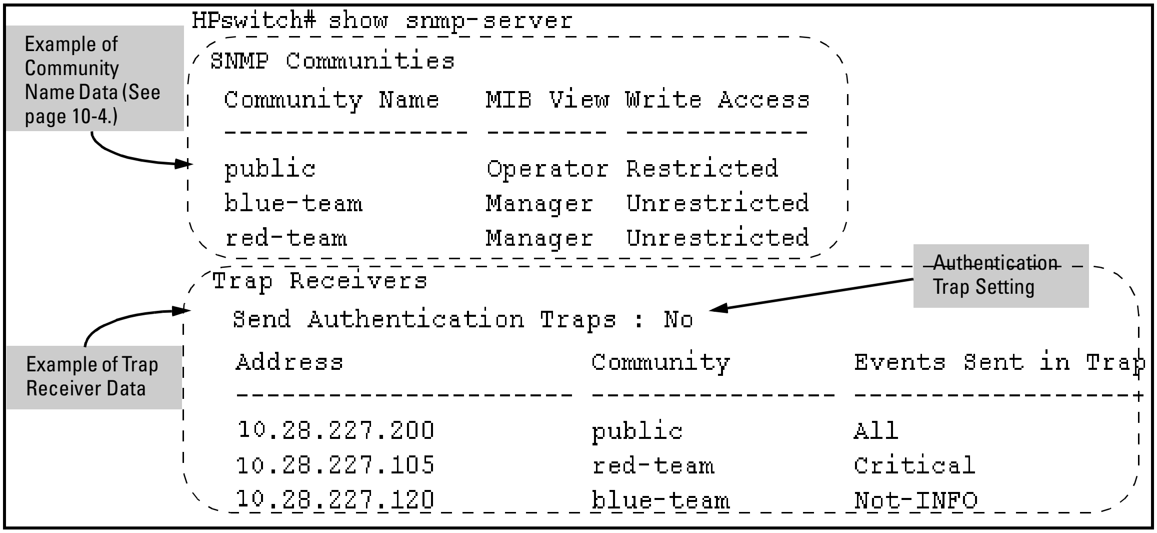

```
HPswitch# show snmp-server
 SNMP Communities
  Community Name   MIB View Write Access
  ---------------- -------- ------------
  public           Operator Restricted
  blue-team        Manager  Unrestricted
  red-team         Manager  Unrestricted
 Trap Receivers
  Send Authentication Traps : No
  Address                 Community        Events Sent in Trap
  ----------------------- ---------------- -------------------
  10.28.227.200           public           All
  10.28.227.105           red-team         Critical
  10.28.227.120           blue-team        Not-INFO
```

Example of Community Name Data (See page 10-4.)

Authentication Trap Setting

Example of Trap Receiver Data

Figure 10-4. Example of Show SNMP-Server Listing

Configuring Trap Receivers. This command specifies trap receivers by community membership, management station IP address, and the type of Event Log messages to send to the trap receiver.

Note

If you specify a community name that does not exist—that is, has not yet been configured on the switch—the switch still accepts the trap receiver assignment. However, no traps will be sent to that trap receiver until the community to which it belongs has been configured on the switch.

Syntax: snmp-server host < *community-string* > < *ip-address* >

Using community name and destination IP address, this command designates a destination network-management station for receiving SNMP event log messages from the switch. If you do not specify the event level, then the switch does not send event log messages as traps. You can specify up to 10 trap receivers (network management stations).

Note: *In all cases, the switch sends any threshold trap(s) to the network management station(s) that explicitly set the threshold(s).*

[<none | all | non-info | critical | debug>]

Options for sending switch Event Log messages to a trap receiver. Refer toTable 10-1, "Options for Sending Event Log Messages as Traps," on page 10-10. The levels specified with these options apply only to Event Log messages, and not to threshold traps.

Table 10-1. Options for Sending Event Log Messages as Traps

Event Level	Description
None (default)	Send no log messages.
All	Send all log messages.
Not INFO	Send the log messages that are not information-only.
Critical	Send critical-level log messages.
Debug	Reserved for HP-internal use.

For example, to configure a trap receiver in a community named "red-team" with an IP address of 10.28.227.130 to receive only "critical" log messages:

```
HPswitch(config)# snmp-server trap-receiver red-team
                 10.28.227.130 critical
```

Notes

To replace one community name with another for the same IP address, you must use **no snmp-server host < community-name> < ip-address >** to delete the unwanted community name. Otherwise, adding a new community name with an IP address already in use with another community name simply creates two allowable community name entries for the same management station.

If you do not specify the event level ([<none | all | non-info | critical | debug>]) then the switch does not send event log messages as traps. "Well-Known" traps and threshold traps (if configured) will still be sent.

Using the CLI To Enable Authentication Traps

Note

For this feature to operate, one or more trap receivers must be configured on the switch. See "In the default configuration, there are no trap receivers configured, and the authentication trap feature is disabled. From the CLI you can configure up to ten SNMP trap receivers to receive SNMP traps from the switch. As an option, you can also configure the switch to send Event Log messages as traps. CLI: Configuring and Displaying Trap Receivers" on page 10-8.

Using the CLI To Enable Authentication Traps.

Syntax: [no] snmp-server enable traps authentication

Enables or disables sending an authentication trap to the configured trap receiver(s) if an unauthorized management station attempts to access the switch.

For example:

```
HPswitch(config)# snmp-server enable traps authentication
```

Check the Event Log in the console interface to help determine why the authentication trap was sent. (Refer to "Using the Event Log To Identify Problem Sources" on page C-22.)

Advanced Management: RMON

The switch supports RMON (Remote Monitoring) on all connected network segments. This allows for troubleshooting and optimizing your network.

The following RMON groups are supported:

- Ethernet Statistics (except the numbers of packets of different frame sizes)
- Alarm
- History (of the supported Ethernet statistics)
- Event

The RMON agent automatically runs in the switch. Use the RMON management station on your network to enable or disable specific RMON traps and events. Note that you can access the Ethernet statistics, Alarm, and Event groups from the HP TopTools for Hubs & Switches network management software. For more on TopTools, see the “Read Me First” document shipped with your switch and also available on HP’s ProCurve web site at

http://www.hp.com/go/hpprocurve

CDP

CDP Features

Feature	Default	Menu	CLI	Web
view the switch's CDP configuration	n/a	—	page 10-20	—
view the switch's CDP Neighbors table	n/a	—	page 10-20	—
clear (reset) the CDP Neighbors table	n/a	—	page 10-21	—
enable or disable CDP on the switch	enabled	—	page 10-22	—
enable or disable CDP operation on an individual port	enabled	—	page 10-23	—
change the transmit interval for the switch's CDP packets	60 seconds	—	page 10-24	—
change the hold time (time-to-live for CDP packets the switch generates)	180 seconds	—	page 10-24	—

Introduction

In Series 5300XL switches, CDP-v1 (Cisco Discovery Protocol, version 1) provides data that aids SNMP-based network mapping utilities designed to discover devices running CDP in a network. To make this data available, the switch transmits information about itself via CDP packets to adjacent devices, and also receives and stores information about adjacent devices running CDP. This enables each CDP device to receive and maintain identity data on each of its CDP neighbors and pass this information off to an SNMP utility designed to query the CDP area of the device's MIB.

Note

To take advantage of CDP in Series 5300XL switches, you should have a working knowledge of SNMP operation and an SNMP utility capable of polling the switches for CDP data. HP's implementation of CDP places specific data into the switch's Management Information Base (MIB). However, retrieval of this data for network mapping is dependent on the operation of your SNMP utility. Refer to the documentation provided with the utility.

An SNMP utility can progressively discover CDP devices in a network by:

1. Reading a given device's CDP Neighbor table (in the Management Information Base, or MIB) to learn about other, neighbor CDP devices
2. Using the information learned in step 1 to go to and read the neighbor devices' CDP Neighbors tables to learn about additional CDP devices, and so on

This section describes CDP operation in a Series 5300XL switches. For information on how to use an SNMP utility to retrieve the CDP information from the switch's CDP Neighbors table (in the switch's MIB), refer to the documentation provided with the particular SNMP utility. For information on the object identifiers in the CDP MIB, see "CDP Neighbor Data and MIB Objects" on page 10-26.

CDP Terminology

- **CDP Device:** A switch, server, router, workstation, or other device running CDP.
- **CDP-Aware:** A device that has CDP in its operating code (with CDP either enabled or disabled in that device).
- **CDP-Disabled**: A CDP-aware device on which CDP is currently disabled.
- **Non-CDP Device:** A device that does not have CDP in its operating code.
- **CDP Neighbor:** A CDP device that is either directly connected to another CDP device or connected to that device by a non-CDP device, such as some hubs.

General CDP Operation

The switch stores information about adjacent CDP devices in a *CDP Neighbors table* maintained in the switch's MIB (Management Information Base). This data is available to SNMP-based applications designed to read CDP data from the MIB. For example:

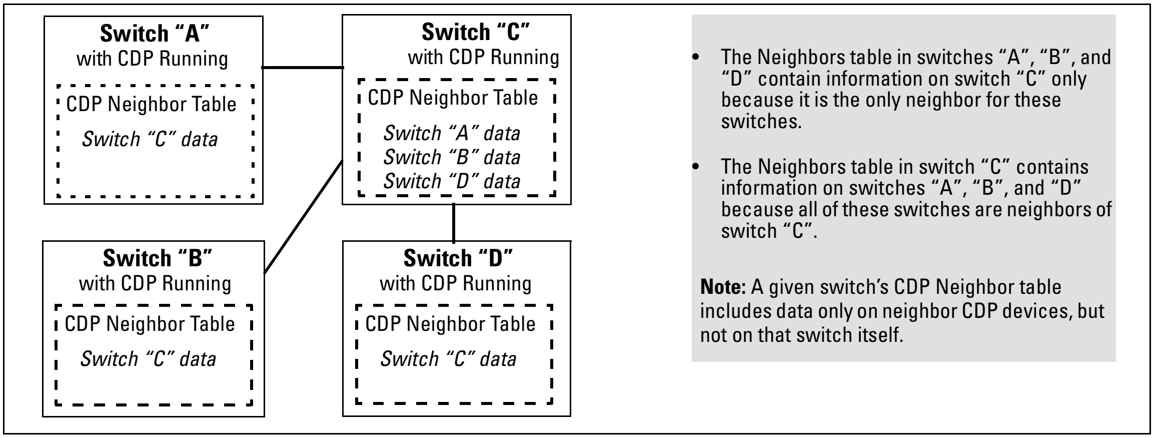

- The Neighbors table in switches "A", "B", and "D" contain information on switch "C" only because it is the only neighbor for these switches.
- The Neighbors table in switch "C" contains information on switches "A", "B", and "D" because all of these switches are neighbors of switch "C".

Note: A given switch's CDP Neighbor table includes data only on neighbor CDP devices, but not on that switch itself.

Figure 10-5. Example of How the Switch Stores Data on Neighbor CDP Devices

Outgoing Packets

A Series 5300XL switch running CDP periodically transmits a one-hop CDP packet out each of its ports. This packet contains data describing the switch and, if the one-hop destination is another device running CDP, the receiving device stores the sending device's data in a CDP Neighbors table. The receiving device also transmits a similar one-hop CDP packet out each of its ports to make itself known to other CDP devices to which it is connected. Thus, each CDP device in the network provides data on itself to the CDP neighbors to which it is directly connected. However, there are instances where a packet is forwarded beyond the immediate neighbor, or simply dropped.

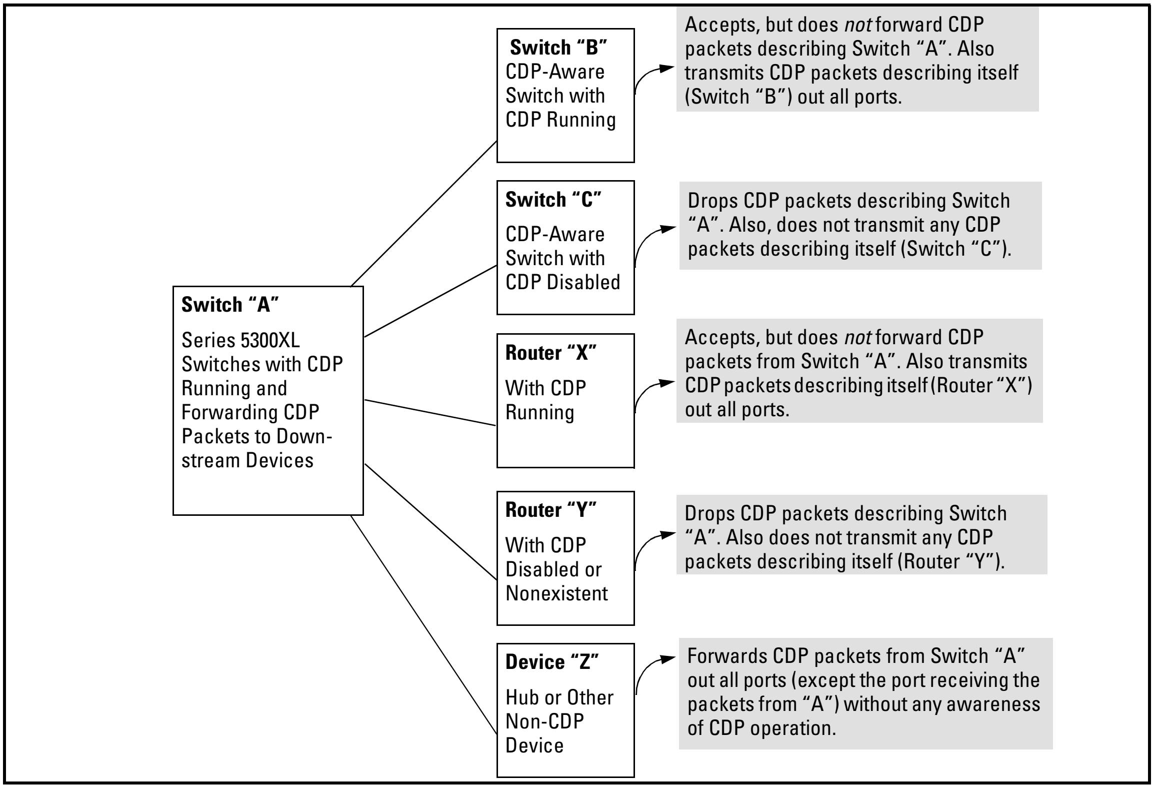
Switch "A"
Series 5300XL Switches with CDP Running and Forwarding CDP Packets to Down-stream Devices
Switch "B"
CDP-Aware Switch with CDP Running
Accepts, but does *not* forward CDP packets describing Switch "A". Also transmits CDP packets describing itself (Switch "B") out all ports.
Switch "C"
CDP-Aware Switch with CDP Disabled
Drops CDP packets describing Switch "A". Also, does not transmit any CDP packets describing itself (Switch "C").
Router "X"
With CDP Running
Accepts, but does *not* forward CDP packets from Switch "A". Also transmits CDP packets describing itself (Router "X") out all ports.
Router "Y"
With CDP Disabled or Nonexistent
Drops CDP packets describing Switch "A". Also does not transmit any CDP packets describing itself (Router "Y").
Device "Z"
Hub or Other Non-CDP Device
Forwards CDP packets from Switch "A" out all ports (except the port receiving the packets from "A") without any awareness of CDP operation.

Figure 10-6. Example of Outgoing CDP Packet Operation

Incoming CDP Packets

When a CDP-enabled Series 5300XL switches receives a CDP packet from another CDP device, it enters that device's data in the CDP Neighbors table, along with the port number where the data was received (and does not forward the packet). The switch also periodically purges the table of any entries that have expired. (The hold time for any data entry in the switch's CDP Neighbors table is configured in the device transmitting the CDP packet, and cannot be controlled in the switch receiving the packet.) A Series 5300XL switch reviews the list of CDP neighbor entries every three seconds, and purges any expired entries.

Non-CDP devices such as some hubs and other devices that do not have CDP capability are transparent to CDP operation. (Other hubs are CDP-aware, but still forward CDP packets as if they were transparent to CDP operation. See "CDP-Capable Hubs" on page 10-29.) However, an intervening CDP-aware device that is CDP-disabled is *not* transparent. For example, in figure 10-7, the

CDP neighbor pairs are as follows: A/1, A/2, A/3, A/B, B/C. Note that “C” and “E” are *not* neighbors because the intervening CDP-disabled switch “D” does not forward CDP packets; i.e. is not transparent to CDP traffic. (For the same reason, switch “E” does not have any CDP neighbors.)

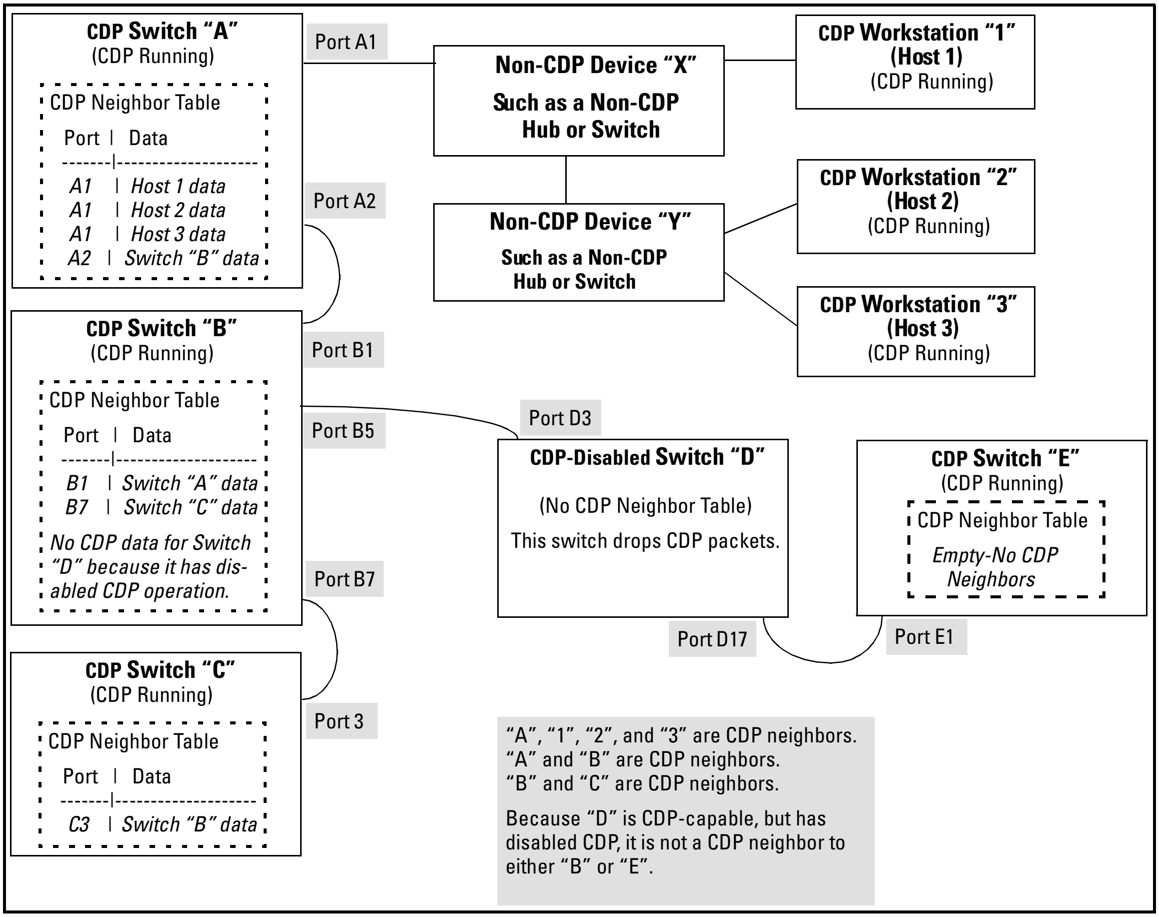

Figure 10-7. Example of Incoming CDP Packet Results

Using the example in figure 10-7, the CDP Neighbor table for switches “A” and “B” would appear similar to these:

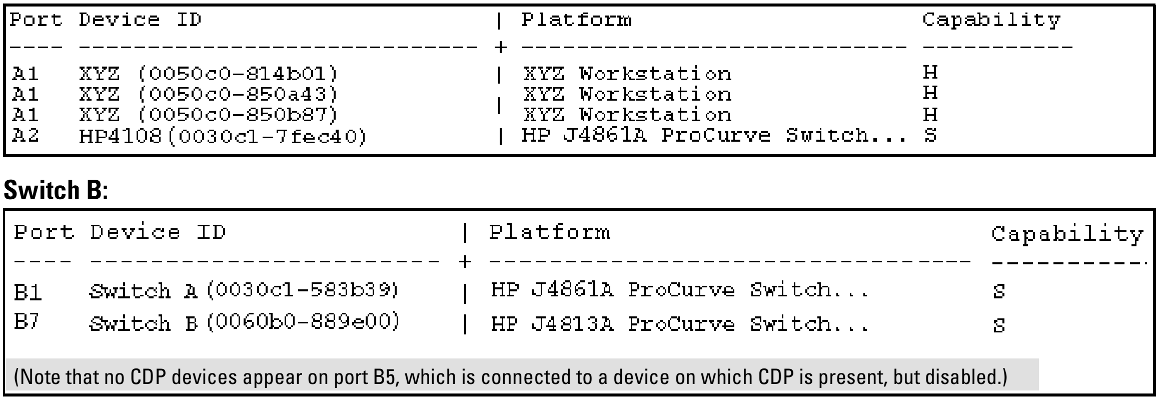

Switch A:

```
Port Device ID                       | Platform                      Capability
---- ------------------------------ + ---------------------------- -----------
A1   XYZ (0050c0-814b01)            | XYZ Workstation               H
A1   XYZ (0050c0-850a43)            | XYZ Workstation               H
A1   XYZ (0050c0-850b87)              XYZ Workstation               H
A2   HP4108(0030c1-7fec40)          | HP J4861A ProCurve Switch... S
```

Switch B:

```
Port Device ID                    | Platform                            Capability
---- ------------------------- + ---------------------------------- ----------
B1   Switch A (0030c1-583b39)   | HP J4861A ProCurve Switch...        S
B7   Switch B (0060b0-889e00)   | HP J4813A ProCurve Switch...        S
```

(Note that no CDP devices appear on port B5, which is connected to a device on which CDP is present, but disabled.)

Figure 10-8. Example of Viewable CDP Neighbor Table for Switches “A” and “B in Figure 10-7

Thus, based on the CDP packets it receives, each CDP device maintains a per-port data entry for each of its neighbors that are running CDP, but not for other CDP devices that are accessible only through a CDP neighbor. (See the relationship between switches A, B, and C in figure 10-7.) In other words, a CDP device will have data on its immediate CDP neighbors (including those reached through a device that is transparent to CDP), but not to other CDP devices in the network.

Table 10-2. How Devices Handle Incoming CDP Packets

Status of Device Receiving a CDP Packet	Action of Receiving Device
Running CDP	Stores neighbor data in CDP Neighbor table. Does not forward CDP packet.
CDP Disabled	Drops CDP packet. There is no CDP Neighbor table and no CDP neighbor data is stored.
No CDP Capability	Forwards CDP packet out all ports except the port on which the packet was received.
Router Running CDP	Stores neighbor data in CDP Neighbor table. Does not forward CDP packet.
Router with CDP (1) Disabled or (2) Not CDP-Capable	Drops CDP packet.

Non-CDP devices (that is, devices that are not capable of running CDP) are transparent to CDP operation. However, an intervening CDP-aware device that is CDP-disabled is *not* transparent. For example, in figure 10-7 (page 10-17), “B”, “D”, and “E” are *not* CDP neighbors because “D” (the intervening

CDP-disabled switch) does not forward CDP packets; i.e. is not transparent to CDP traffic. (For the same reason, switch "E" does not have any CDP neighbors.)

Figure 10-7 (page 10-17) illustrates how multiple CDP neighbors can appear on a single port. In this case, switch "A" has three CDP neighbors on port 1 because the intervening devices are not CDP-capable and simply forward CDP neighbors data out all ports (except the port on which the data was received).

Default Configuration. In the factory-default configuration, CDP is enabled and running on all ports. In this case, the **holdtime** is 180 seconds and the **timer** (CDP Transmit Interval) is 60 seconds.

Configuring CDP on the Switch

On a Series 5300XL switch you can:

- View the switch's current global and per-port CDP configuration
- List the current contents of the switch's CDP Neighbors table (that is, view a listing of the CDP devices of which the switch is aware)
- Enable or disable CDP (Default: Enabled)
- Specify the hold time (CDP packet time-to-live) for CDP data delivered to neighboring CDP devices. For example, in CDP switch "A" you can specify the hold time for switch "A" entries in the CDP Neighbor tables of other CDP devices. (Default: 180 seconds)
- Specify the transmission interval for CDP packets. (Default: 60 seconds)

CLI: Viewing and Configuring CDP

CDP Commands	Page
show CDP	10-20
show CDP neighbors	10-20
cdp clear	10-21
[no] cdp run	10-22
[no] cdp enable	10-23
cdp holdtime	10-24
cdp timer	10-24

Viewing the Switch's Current CDP Configuration

Syntax: show cdp

Lists the switch's global and per-port CDP configuration.

This example shows the default CDP configuration.

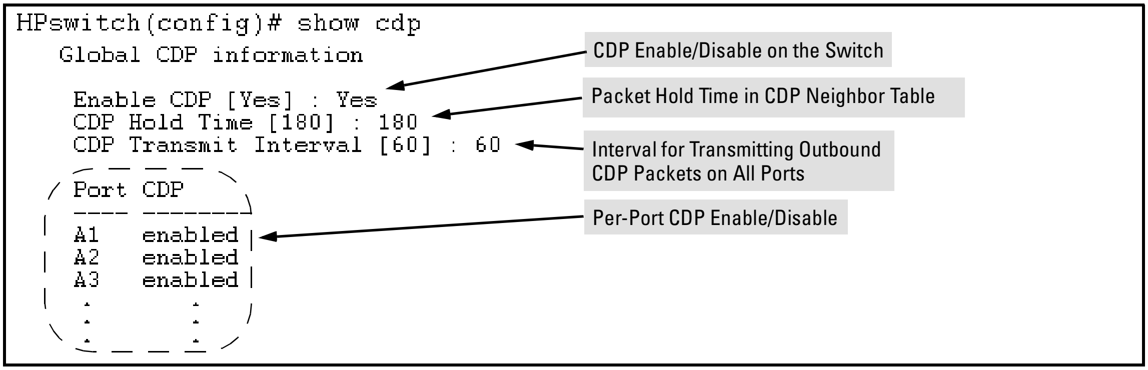

```
HPswitch(config)# show cdp
   Global CDP information

    Enable CDP [Yes] : Yes
    CDP Hold Time [180] : 180
    CDP Transmit Interval [60] : 60

    Port CDP
    ---- --------
    A1   enabled
    A2   enabled
    A3   enabled
     .        .
     .        .
     .        .
```

CDP Enable/Disable on the Switch

Packet Hold Time in CDP Neighbor Table

Interval for Transmitting Outbound CDP Packets on All Ports

Per-Port CDP Enable/Disable

Figure 10-9. Example of Show CDP with the Default CDP Configuration

Viewing the Switch's Current CDP Neighbors Table

Devices are listed by the port on which they were detected.

Syntax: show cdp neighbors

Lists the neighboring CDP devices the switch detects, with a subset of the information collected from the device's CDP packet. (For more on this topic, refer to table 10-3, "CDP Neighbors Data" on page 10-27.)

[detail]

Provides a longer list of details on all of the CDP-aware device the switch detects.

[[e] *port-num*]

Lists the details for the CDP-aware device connected to the specified port. (Allows only one port at a time.)

(For more on this topic, see "CDP Neighbor Data and MIB Objects" on page 10-26.)

Figure 10-10 lists six CDP devices (four switches and two workstations) that the switch has detected by receiving their CDP packets.

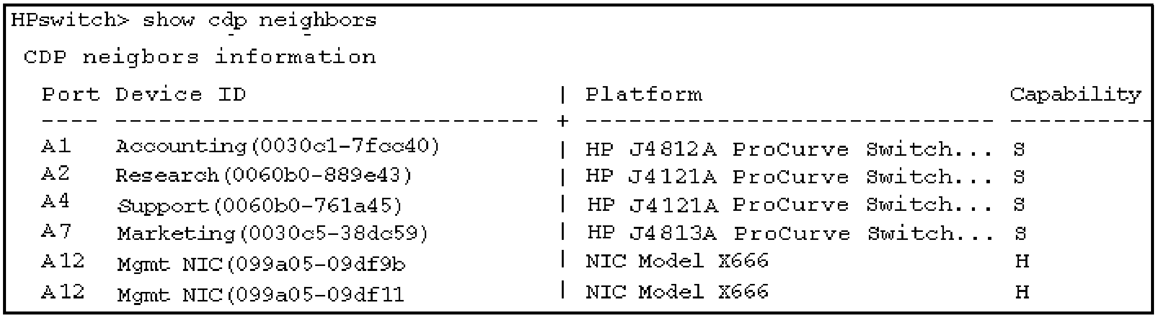

```
HPswitch> show cdp neighbors
 CDP neigbors information
  Port Device ID                      | Platform                     Capability
  ---- ------------------------------ + ---------------------------- ----------
  A1   Accounting(0030c1-7fcc40)      | HP J4812A ProCurve Switch... S
  A2   Research(0060b0-889e43)        | HP J4121A ProCurve Switch... S
  A4   Support(0060b0-761a45)         | HP J4121A ProCurve Switch... S
  A7   Marketing(0030c5-38dc59)       | HP J4813A ProCurve Switch... S
  A12  Mgmt NIC(099a05-09df9b         | NIC Model X666               H
  A12  Mgmt NIC(099a05-09df11         | NIC Model X666               H
```

Figure 10-10. Example of CDP Neighbors Table Listing

Figure 10-11 illustrates a topology of CDP-enabled devices for the CDP Neighbors table listing in figure 10-10.

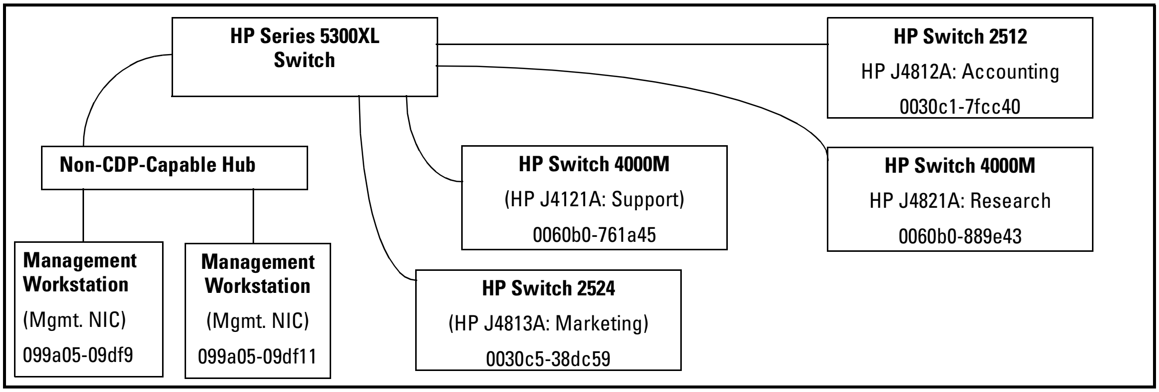

Figure 10-11. Example of CDP-Enabled Devices in a Topology for the Listing in Figure 10-10

Clearing (Resetting) the CDP Neighbors Table

Syntax: cdp clear

Removes any records of CDP neighbor devices from the switch' s CDP MIB objects.

If you execute **cdp clear** and then execute **show cdp** neighbors before the switch receives a CDP packet from any neighbor device, the displayed table appears empty.

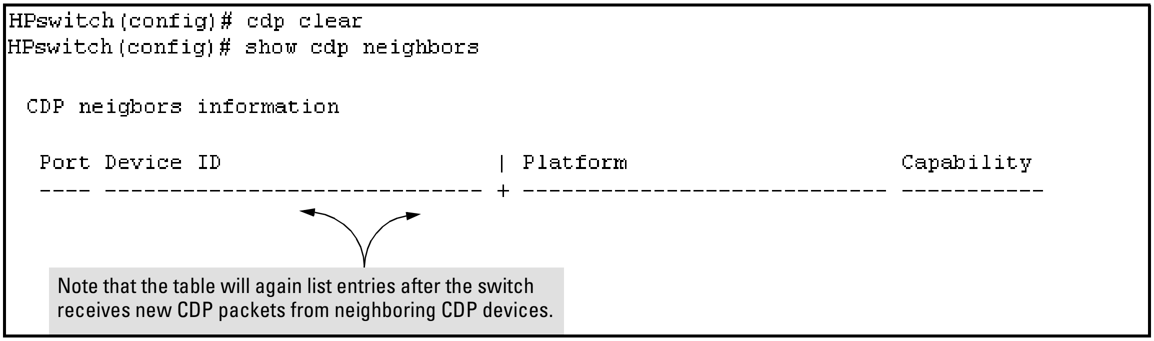

```
HPswitch(config)# cdp clear
HPswitch(config)# show cdp neighbors

 CDP neigbors information

  Port Device ID                      | Platform                     Capability
  ---- ------------------------------ + ---------------------------- -----------
```

Note that the table will again list entries after the switch receives new CDP packets from neighboring CDP devices.

Figure 10-12. View of the CDP Neighbors Table Immediately After Executing cdp clear

Configuring CDP Operation

Enabling or Disabling CDP Operation on the Switch. Enabling CDP operation (the default) on the switch causes the switch to:

- Transmit CDP packets describing itself to other, neighboring CDP devices
- Add entries to its CDP Neighbors table for any CDP packets it receives from other, neighboring CDP devices

Disabling CDP operation clears the switch's CDP Neighbors table, prevents the switch from transmitting outbound CDP packets to advertise itself to neighboring CDP devices, and causes the switch to drop inbound CDP packets from other devices without entering the data in the CDP Neighbors table.

Syntax: [no] cdp run

Enables or disables CDP operation on the switch. (Default: Enabled)

For example, to disable CDP on the switch:

```
HPswitch(config) no cdp run
```

When CDP is disabled:

- **show cdp neighbors** displays an empty CDP Neighbors table
- **show cdp** displays

```
Global CDP information
 Enable CDP [Yes] : No
```

Enabling or Disabling CDP Operation on Individual Ports. In the factory-default configuration, the switch has all ports enabled and transmitting CDP packets. Disabling CDP on a port prevents that port from sending outbound CDP packets and causes it to drop inbound CDP packets without recording their data in the CDP Neighbors table. Suppose, for example, that switches "A" and "B" in figure 10-13 (below) are running CDP, and that port A1 on switch "A" is connected to port B5 on switch "B". If you disable CDP on port A1 of switch "A", then switch "B" will no longer receive CDP packets from switch "A" and switch "A" will drop the CDP packets it receives from switch "B".

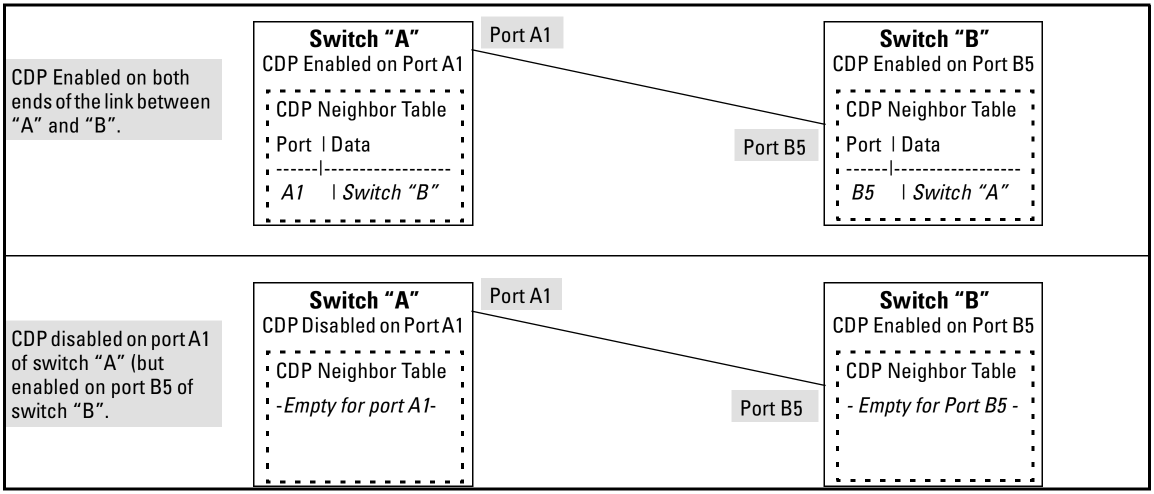

Figure 10-13. Example of Disabling CDP on an Individual Port

(The switch "A" entry in the switch "B" CDP Neighbors table remains until the **cdp holdtime** (time-to-live; set in switch "A") expires. Until then, the **show cdp neighbors** command continues to list switch "A" on port B5 of switch "B".)

Syntax: [no] cdp enable <[ethernet] port-list>

For example, to disable CDP on port A1 of a Series 5300XL switch:

```
HPswitch(config) no cdp enable a1
```

Changing the Transmission Interval for Outbound CDP Packets.

Syntax: cdp timer < 5 . . 254 >

Changes the interval the switch uses to transmit CDP packets describing itself to neighbor devices. (Default: 60 seconds)

For example, if the switch's transmit interval for CDP packets was set to a non-default value, you would use this command to reset it to one minute:

```
HPswitch(config) cdp timer 60
```

Changing the Hold Time (CDP Packet Time-To-Live) for a Switch's CDP Packet Information. This parameter is controlled in the transmitting switch, and applies to all outbound CDP packets the switch transmits.

Syntax: cdp holdtime < 5 . . 254 >

Changes the hold time for the switch's CDP packet information in the CDP Neighbors table of another CDP-aware device. (Default: 180 seconds; Range: 10 - 255 seconds.)

For example, to configure a switch's outbound CDP packets to live for one minute in the CDP Neighbors table of neighboring CDP devices:

```
HPswitch(config) cdp holdtime 60
```

Effect of Spanning Tree (STP) On CDP Packet Transmission

If STP has blocked a port on the switch, that port does not transmit CDP packets. However, the port still receives CDP packets if the device on the other end of the link has CDP enabled. Thus, for example, if switch "A" has two ports linked to switch "B" (which is a CDP neighbor and also the STP root device) and STP blocks traffic on one port and forwards traffic on the other:

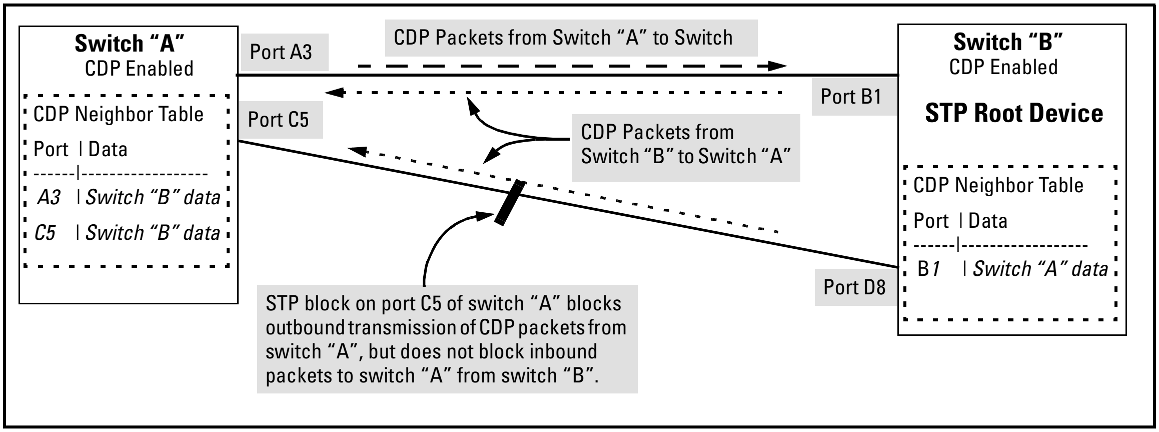

Figure 10-14. Example of STP Effect on CDP Packet Transmission

- Switch "A" sends outbound CDP packets on the forwarding link, and the switch "B" CDP Neighbors table shows switch "A" on only one port.
- Switch "B" sends outbound CDP packets on both links, and the switch "A" CDP Neighbors table shows switch "B" on both ports.

To summarize, in a CDP neighbor pair running STP with redundant links, if one of the switches is the STP root, it transmits CDP packets out all ports connecting the two switches, while the other switch transmits CDP packets out only the unblocked port. Thus, the STP root switch will appear on multiple ports in the non-root switch's CDP Neighbors table, while the non-root switch will appear on only one port in the root switch's CDP Neighbors table.

Selection of the IP Address To Include in Outbound CDP Packets

A switch with CDP enabled uses the following prioritized criteria to determine which IP address to include in its outbound CDP packets:

1. If only one VLAN on the port has an IP address, the switch uses that IP address.
2. If the Primary VLAN on the port has an IP address, the switch uses the Primary VLAN IP address.
3. If 1 and 2 do not apply, then the switch determines which VLANs on the port have IP addresses and uses the IP address of the VLAN with the lowest VID (VLAN Identification number) in this group.

4. If a CDP switch does not detect an IP address on the connecting port of a CDP neighbor, then the loopback IP address is used (127.0.0.1).

For example, in figure 10-15, port A1 on CDP switch “X” is connected to port C5 on CDP neighbor switch “Y”, with the indicated VLAN configuration on port C5:

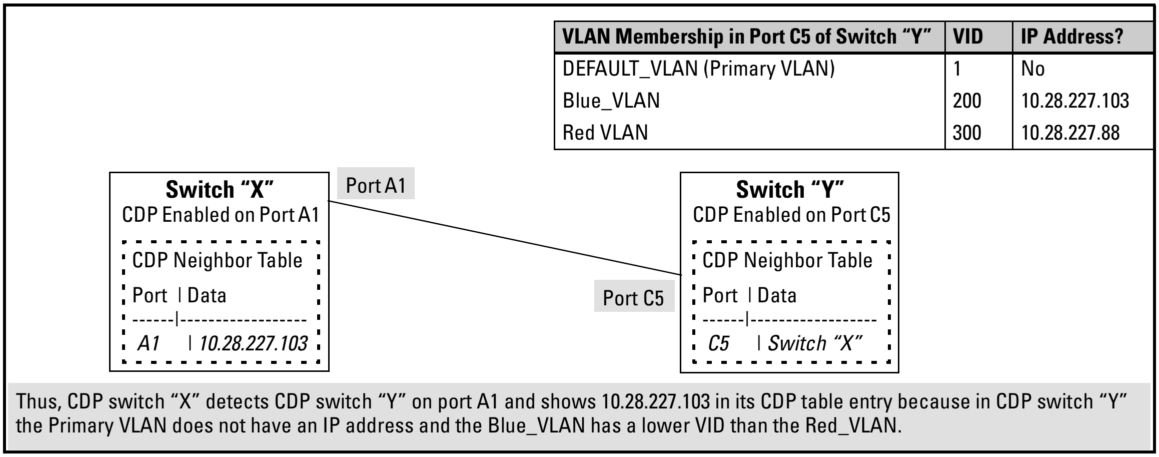

VLAN Membership in Port C5 of Switch “Y”	VID	IP Address?
DEFAULT_VLAN (Primary VLAN)	1	No
Blue_VLAN	200	10.28.227.103
Red VLAN	300	10.28.227.88

Thus, CDP switch “X” detects CDP switch “Y” on port A1 and shows 10.28.227.103 in its CDP table entry because in CDP switch “Y” the Primary VLAN does not have an IP address and the Blue_VLAN has a lower VID than the Red_VLAN.

Figure 10-15. Example of IP Address Selection when a CDP Neighbor Has Multiple VLANs with IP Addresses

CDP Neighbor Data and MIB Objects

The switch places the data received from inbound CDP packets into its MIB (Management Information Base). This data is available in three ways:

- Using the switch’s **show cdp neighbors** command to display a subset of Neighbor data
- Using the **walkmib** command to display a listing of the CDP MIB objects
- Electronically, using an SNMP utility designed to search the MIB for CDP data

As shown under “Viewing the Switch’s Current CDP Neighbors Table” on page 10-20, you can list a subset of data for each CDP device currently found in the switch’s CDP Neighbors table. Table 10-3, “CDP Neighbors Data”, describes the CDP Neighbor data set available in the Series 5300XL switches.

Table 10-3. CDP Neighbors Data

CDP Neighbor Data	Displayed Neighbors Table	MIB	
Address Type	No	Yes	Always "1" (IP address only).
CDP Cache Address	No	Yes	IP address of source device.
Software Version	Yes	Yes	ASCII String
Device Name (ASCII string)	Yes	Yes	In HP Procurve switches, this is the value configured for the System Name parameter.
Device MAC Address	Yes	Yes	Included in the Device Name entry.
Destination Port Number	Yes	Yes	On the Series 5300XL switches (the receiving device), the number of the port through which the CDP packet arrived.
Source Port Number	No	Yes	On the source (neighbor) device, the number of the port through which the CDP packet was sent.
Product Name (ASCII string)	Yes	Yes	Platform name designated by vendor.
Capability Code (Device Type)	Yes (alpha character)	Yes (numeric character)	1 or R: Router 2: Transparent Bridge 4 or B: Source Route Bridge 8 or S: Switch 16 or H: Host 32 or I: IGMP conditional filtering 64 or r: Repeater

Displaying CDP Neighbor Data.

Syntax: walkmib CdpCacheEntry

Displays the superset of CDP neighbor held in the MIB.

For example, with two CDP devices connected to ports A1 and A3 on the switch, you would see a **walkmib** listing similar to this:

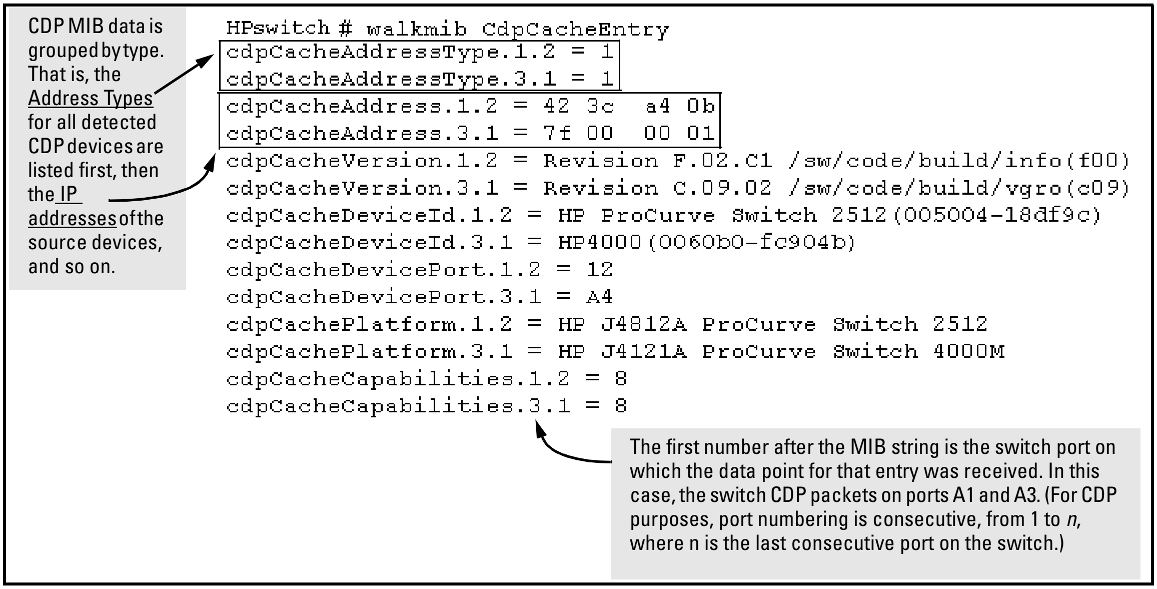

CDP MIB data is grouped by type. That is, the Address Types for all detected CDP devices are listed first, then the IP addresses of the source devices, and so on.

```
HPswitch # walkmib CdpCacheEntry
cdpCacheAddressType.1.2 = 1
cdpCacheAddressType.3.1 = 1
cdpCacheAddress.1.2 = 42 3c  a4 0b
cdpCacheAddress.3.1 = 7f 00  00 01
cdpCacheVersion.1.2 = Revision F.02.C1 /sw/code/build/info(f00)
cdpCacheVersion.3.1 = Revision C.09.02 /sw/code/build/vgro(c09)
cdpCacheDeviceId.1.2 = HP ProCurve Switch 2512(005004-18df9c)
cdpCacheDeviceId.3.1 = HP4000(0060b0-fc904b)
cdpCacheDevicePort.1.2 = 12
cdpCacheDevicePort.3.1 = A4
cdpCachePlatform.1.2 = HP J4812A ProCurve Switch 2512
cdpCachePlatform.3.1 = HP J4121A ProCurve Switch 4000M
cdpCacheCapabilities.1.2 = 8
cdpCacheCapabilities.3.1 = 8
```

The first number after the MIB string is the switch port on which the data point for that entry was received. In this case, the switch CDP packets on ports A1 and A3. (For CDP purposes, port numbering is consecutive, from 1 to *n*, where n is the last consecutive port on the switch.)

Figure 10-16. Example of CDP Neighbor Data in a Series 5300XL Switch MIB

For the current Series 5300XL switch MIB, go the HP Procurve World Wide Web site at:

http://www.hp.com/go/hpprocurve

Click on **software**, then **MIBs**.

Operating Notes

Neighbor Maximum. The Series 5300XL switches support up to 60 entries (neighbors) in the CDP Neighbors table. Remember that multiple CDP devices can be neighbors on the same port if they are connected to the switch through a non-CDP device.

CDP Version Data. The Series 5300XL switches use CDP-V1, but do not include IP prefix information, which is a router function; not a switch application.

Port Trunking with CDP. Where a static or LACP trunk forms the link between the switch and another CDP device, only one physical link in the trunk is used to transmit outbound CDP packets.

CDP-Capable Hubs. Some hubs are capable of running CDP, but also forward CDP packets as if the hub itself were transparent to CDP. Such hubs will appear in the switch's CDP Neighbor table and will also maintain a CDP neighbor table similar to that for switches. For more information, refer to the documentation provided for the specific hub.

Troubleshooting CDP Operation. Turn to "Using the Event Log To Identify Problem Sources" on page C-22.

11

Port-Based Virtual LANs (VLANs) and GVRP

Contents

Overview

This chapter describes the following features and how to configure them with the switch's built-in interfaces:

- **Port-Based VLANs — Page 11-3:**
- **GVRP — Page 11-34:**

For general information on how to use the switch's built-in interfaces, see:

- Chapter 2, "Using the Menu Interface"
- Chapter 3, "Using the Command Line Interface (CLI)"
- Chapter 4, "Using the HP Web Browser Interface
- Chapter 5, "Switch Memory and Configuration"

Port-Based Virtual LANs (Static VLANs)

VLAN Features

Feature	Default	Menu	CLI	Web
view existing VLANs	n/a	page 11-10 thru 11-15	page 11-16	page 11-21
configuring static VLANs	default VLAN with VID = 1	page 11-10 thru 11-15	page 11-15	page 11-21
configuring dynamic VLANs	disabled	See "GVRP" on page 11-34.		

A VLAN is a group of ports designated by the switch as belonging to the same broadcast domain. (That is, all ports carrying traffic for a particular subnet address would normally belong to the same VLAN.)

Note

This section describes *static* VLANs, which are VLANs you manually configure with a name, VLAN ID (VID), and port assignments. (For information on *dynamic* VLANs, see "GVRP" on page 11-34.)

Using a VLAN, you can group users by logical function instead of physical location. This helps to control bandwidth usage by allowing you to group high-bandwidth users on low-traffic segments and to organize users from different LAN segments according to their need for common resources.

By default, the Series 5300XL switch is 802.1Q VLAN enabled and allow up to 256 port-based VLANs (default: 8). For information on GVRP, see "GVRP" on page 11-34. (The 802.1Q compatibility enables you to assign each switch port to multiple VLANs, if needed, and the port-based nature of the configuration allows interoperation with older switches that require a separate port for each VLAN.)

General Use and Operation. Port-based VLANs are typically used to enable broadcast traffic reduction and to increase security. A group of network users assigned to a VLAN form a broadcast domain that is separate from other VLANs that may be configured on a switch. On a given switch, packets are forwarded only between ports that are designated for the same VLAN. Thus, all ports carrying traffic for a particular subnet address should be configured to the same VLAN. Cross-domain broadcast traffic in the switch is eliminated and bandwidth is saved by not allowing packets to flood out all ports. An external router is required to enable separate VLANs on a switch to communicate with each other.

For example, referring to figure 11-1, if ports A1 through A4 belong to VLAN_1 and ports A5 through A8 belong to VLAN_2, traffic from end-node stations on ports A2 through A4 is restricted to only VLAN_1, while traffic from ports A5 through A7 is restricted to only VLAN_2. For nodes on VLAN_1 to communicate with VLAN_2, their traffic must go through an external router via ports A1 and A8.

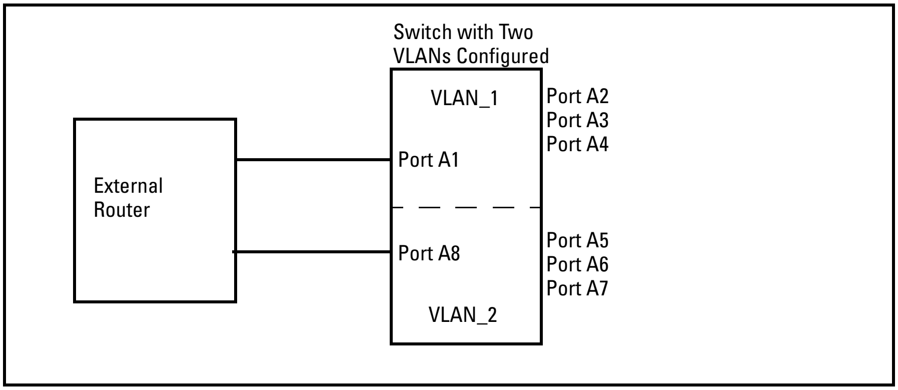

Figure 11-1. Example of Routing Between VLANs via an External Router

Overlapping (Tagged) VLANs. A port on the Series 5300XL switches can be a member of more than one VLAN if the device to which they are connected complies with the 802.1Q VLAN standard. For example, a port connected to a central server using a network interface card (NIC) that complies with the 802.1Q standard can be a member of multiple VLANs, allowing members of multiple VLANs to use the server. Although these VLANs cannot communicate with each other through the server, they can all access the server *over the same connection from the switch.* Where VLANs overlap in this way, VLAN "tags" are used to distinguish between traffic from different VLANs.

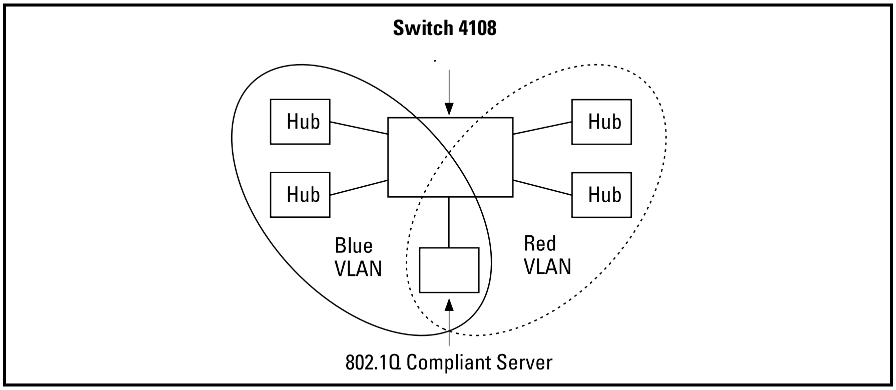

Figure 11-2. Example of Overlapping VLANs Using the Same Server

Similarly, using 802.1Q-compliant switches, you can connect multiple VLANs through a single switch-to-switch link.

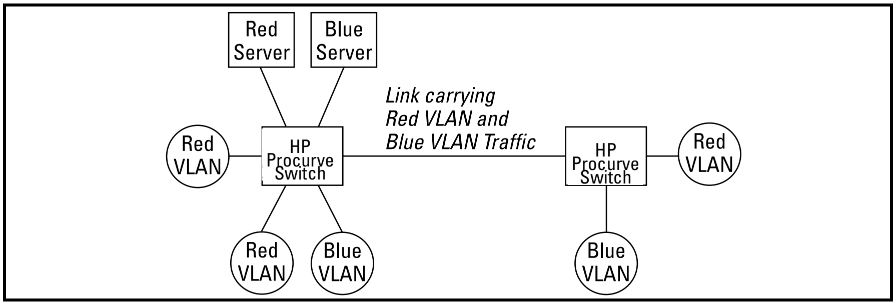

Figure 11-3. Example of Connecting Multiple VLANs Through the Same Link

Introducing Tagged VLAN Technology into Networks Running Legacy (Untagged) VLANs. You can introduce 802.1Q-compliant devices into networks that have built untagged VLANs based on earlier VLAN technology. The fundamental rule is that legacy/untagged VLANs require a separate link for each VLAN, while 802.1Q, or tagged VLANs can combine several VLANs in one link. This means that on the 802.1Q-compliant device, separate ports (configured as untagged) must be used to connect separate VLANs to non-802.1Q devices.

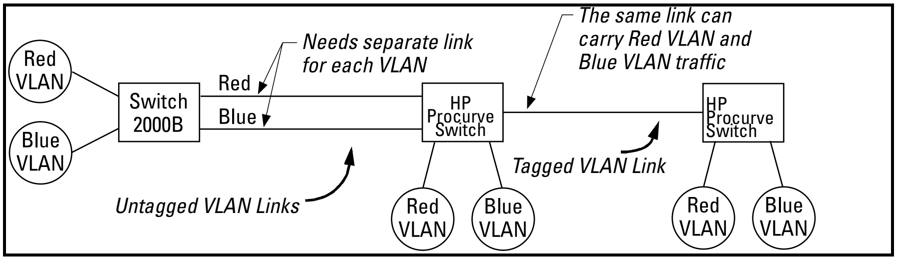

Figure 11-4. Example of Tagged and Untagged VLAN Technology in the Same Network

For more information on VLANs, refer to:

- "Overview of Using VLANs" (page 11-6)
- "Menu: Configuring VLAN Parameters (page 11-10)
- "CLI: Configuring VLAN Parameters" (page 11-10)
- "Web: Viewing and Configuring VLAN Parameters" (page 11-21)
- "VLAN Tagging Information" (page 11-22)
- "Effect of VLANs on Other Switch Features" (page 11-31)
- "VLAN Restrictions" (page 11-32)

Overview of Using VLANs

VLAN Support and the Default VLAN

In the factory default configuration, VLAN support is enabled and all ports on the switch belong to the default VLAN (named DEFAULT_VLAN). This places all ports in the switch into one physical broadcast domain. In the factory-default state, the default VLAN is the primary VLAN.

You can partition the switch into multiple virtual broadcast domains by adding one or more additional VLANs and moving ports from the default VLAN to the new VLANs. (The switch supports up to 256 VLANs.) You can change the name of the default VLAN, but you cannot change the default VLAN's VID (which is always "1"). Although you can remove all ports from the default VLAN, this VLAN is always present; that is, you cannot delete it from the switch.

The Primary VLAN

Because certain features and management functions run on only one VLAN in the switch, and because DHCP and Bootp can run per-VLAN, there is a need for a dedicated VLAN to manage these features and ensure that multiple instances of DHCP or Bootp on different VLANs do not result in conflicting configuration values for the switch. The *primary* VLAN is the VLAN the switch uses to run and manage these features and data. In the factory-default configuration, the switch designates the default VLAN (DEFAULT_VLAN) as the primary VLAN. However, to provide more control in your network, you can designate another VLAN as primary. To summarize, *designating a non-default VLAN as primary* means that:

- The switch reads DHCP responses on the primary VLAN instead of on the default VLAN. (This includes such DHCP-resolved parameters as the TimeP server address, Default TTL, and IP addressing—including the Gateway IP address—when the switch configuration specifies DHCP as the source for these values.)
- The default VLAN continues to operate as a standard VLAN (except, as noted above, you cannot delete it or change its VID).
- Any ports not specifically assigned to another VLAN will remain assigned to the Default VLAN, regardless of whether it is the primary VLAN.

Candidates for primary VLAN include any static VLAN currently configured on the switch. (A dynamic—GVRP-learned—VLAN that has not been converted to a static VLAN cannot be the primary VLAN.) To display the current primary VLAN, use the CLI **show vlan** command.

Note

If you configure a non-default VLAN as the primary VLAN, you cannot delete that VLAN unless you first select a different VLAN to act as primary.

If you manually configure a gateway on the switch, it will ignore any gateway address received via DHCP or Bootp.

Per-Port Static VLAN Configuration Options

The following figure and table show the options you have for assigning individual ports to a static VLAN. Note that GVRP, if configured, affects these options and VLAN behavior on the switch. The display below shows the per-port VLAN configuration options. Table 11-1 briefly describes these options.

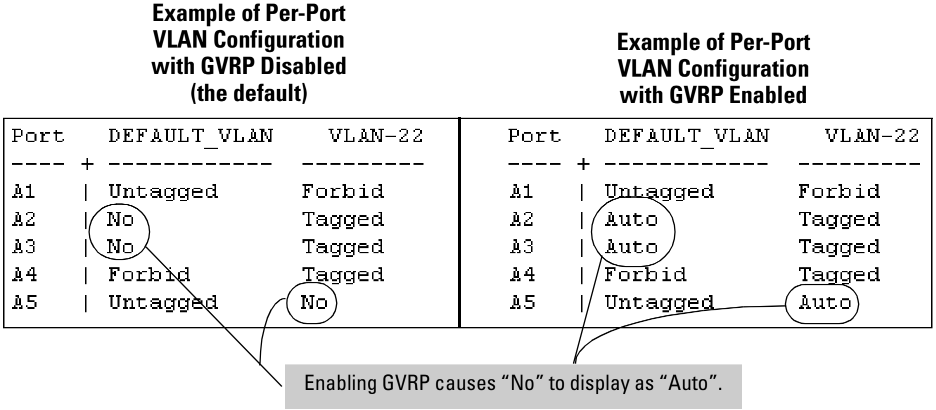

Example of Per-Port VLAN Configuration with GVRP Disabled (the default)

```
Port   DEFAULT_VLAN    VLAN-22
---- + ------------  ---------
A1   | Untagged       Forbid
A2   | No             Tagged
A3   | No             Tagged
A4   | Forbid         Tagged
A5   | Untagged       No
```

Example of Per-Port VLAN Configuration with GVRP Enabled

```
Port   DEFAULT_VLAN    VLAN-22
---- + ------------  ---------
A1   | Untagged       Forbid
A2   | Auto           Tagged
A3   | Auto           Tagged
A4   | Forbid         Tagged
A5   | Untagged       Auto
```

Enabling GVRP causes “No” to display as “Auto”.

Figure 11-5. Comparing Per-Port VLAN Options With and Without GVRP

Table 14-1. Per-Port VLAN Configuration Options

Parameter	Effect on Port Participation in Designated VLAN
Tagged	Allows the port to join multiple VLANs.
Untagged	Allows VLAN connection to a device that is configured for an untagged VLAN instead of a tagged VLAN. The switch allows no more than one untagged VLAN assignment per port.
No *- or -* **Auto**	**No**: Appears when the switch is not GVRP-enabled; prevents the port from joining that VLAN. **Auto**: Appears when GVRP is enabled on the switch; allows the port to dynamically join any advertised VLAN that has the same VID
Forbid	Prevents the port from joining the VLAN, regardless of whether GVRP is enabled on the switch.

General Steps for Using VLANs

1. Plan your VLAN strategy and create a map of the logical topology that will result from configuring VLANs. Include consideration for the interaction between VLANs and other features such as Spanning Tree Protocol, load balancing, and IGMP. (Refer to "Effect of VLANs on Other Switch Features" on page 11-31.) If you plan on using dynamic VLANs, include the port configuration planning necessary to support this feature. (See "GVRP" on page 11-34.)

   By default, VLAN support is enabled and the switch is configured for eight VLANs.

2. Configure at least one VLAN in addition to the default VLAN.

3. Assign the desired switch ports to the new VLAN(s).

4. If you are managing VLANs with SNMP in an IP network, each VLAN must have an IP address. Refer to "IP Configuration" on page 7-3.

Notes on Using VLANs

- If you are using DHCP/Bootp to acquire the switch's configuration, packet time-to-live, and TimeP information, you must designate the VLAN on which DHCP is configured for this purpose as the primary VLAN. (In the factory-default configuration, the DEFAULT_VLAN is the primary VLAN.)
- IGMP, and some other features operate on a "per VLAN" basis. This means you must configure such features separately for each VLAN in which you want them to operate.
- You can rename the default VLAN, but you cannot change its VID (1) or delete it from the switch.
- Any ports *not* specifically assigned to another VLAN will remain assigned to the DEFAULT_VLAN.
- To delete a VLAN from the switch, you must first remove from that VLAN any ports assigned to it.
- Changing the number of VLANs supported on the switch requires a reboot. Other VLAN configuration changes are dynamic.

Menu: Configuring VLAN Parameters

In the factory default state, support is enabled for up to eight VLANs. (You can change the switch VLAN configuration to support up to 256 VLANs.) Also, all ports on the switch belong to the default VLAN (DEFAULT_VLAN) and are in the same broadcast/multicast domain. (The default VLAN is also the default primary VLAN—see "The Primary VLAN" on page 11-6.) In addition to the default VLAN, you can configure up to 255 other static VLANs by changing the "Maximum VLANs" parameter, adding new VLAN names and VIDs, and then assigning one or more ports to each VLAN. (The switch accepts a maximum of 256 VLANs, including the default VLAN and any dynamic VLANs the switch creates if you enable GVRP—page 11-34.) Note that each port can be assigned to multiple VLANs by using VLAN tagging. (See "VLAN Tagging Information" on page 11-22.)

To Change VLAN Support Settings

This section describes:

- Changing the maximum number of VLANs to support
- Changing the primary VLAN selection (See "Changing the Primary VLAN" on page 11-18.)
- Enabling or disabling dynamic VLANs (See "GVRP" on page 11-34.)

1. From the Main Menu select:

   2. Switch Configuration

      8. VLAN Menu . . .

         1. VLAN Support

   You will then see the following screen:

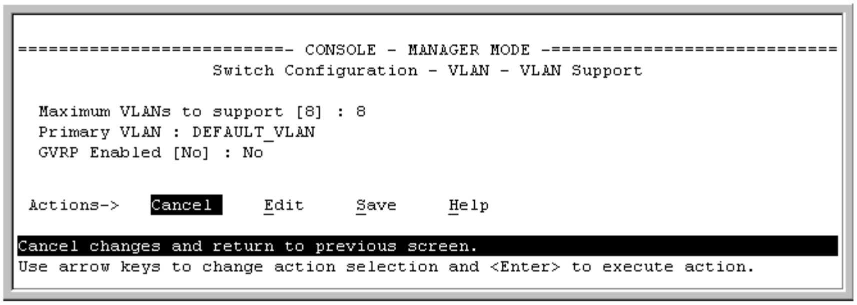

```
==========================- CONSOLE - MANAGER MODE -============================
                      Switch Configuration - VLAN - VLAN Support

  Maximum VLANs to support [8] : 8
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled [No] : No

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 11-6. The Default VLAN Support Screen

2. Press E (for **Edit**), then do one or more of the following:

- To change the maximum number of VLANs, type the new number (1 - 256 allowed; default 8).
- To designate a different VLAN as the primary VLAN, select the **Primary VLAN** field and use the space bar to select from the existing options.
- To enable or disable dynamic VLANs, select the **GVRP Enabled** field and use the Space bar to toggle between options. (For GVRP information, see "GVRP" on page 11-34.)

Note

For optimal switch memory utilization, set the number of VLANs at the number you will likely be using or a few more. If you need more VLANs later, you can increase this number, but a switch reboot will be required at that time.

3. Press [Enter] and then [S] to save the VLAN support configuration and return to the VLAN Menu screen.

   If you changed the value for **Maximum VLANs to support**, you will see an asterisk next to the **VLAN Support** option (see below).

An asterisk indicates you must reboot the switch to implement the new Maximum VLANs setting.

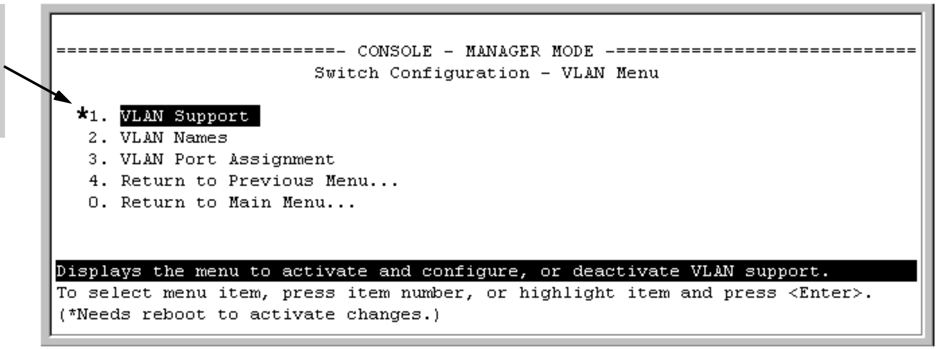

```
==========================- CONSOLE - MANAGER MODE -==========================
                      Switch Configuration - VLAN Menu

  *1. VLAN Support
   2. VLAN Names
   3. VLAN Port Assignment
   4. Return to Previous Menu...
   0. Return to Main Menu...


Displays the menu to activate and configure, or deactivate VLAN support.
To select menu item, press item number, or highlight item and press <Enter>.
(*Needs reboot to activate changes.)
```

Figure 11-7. VLAN Menu Screen Indicating the Need To Reboot the Switch

   – If you changed the VLAN Support option, you must reboot the switch before the Maximum VLANs change can take effect. You can go on to configure other VLAN parameters first, but remember to reboot the switch when you are finished.
   – If you did not change the VLAN Support option, a reboot is not necessary.

4. Press [0] to return to the Main Menu.

Adding or Editing VLAN Names

Use this procedure to add a new VLAN or to edit the name of an existing VLAN.

1. From the Main Menu select:

   2. Switch Configuration
   8. VLAN Menu . . .
   2. VLAN Names

   If multiple VLANs are not yet configured you will see a screen similar to figure 11-8:

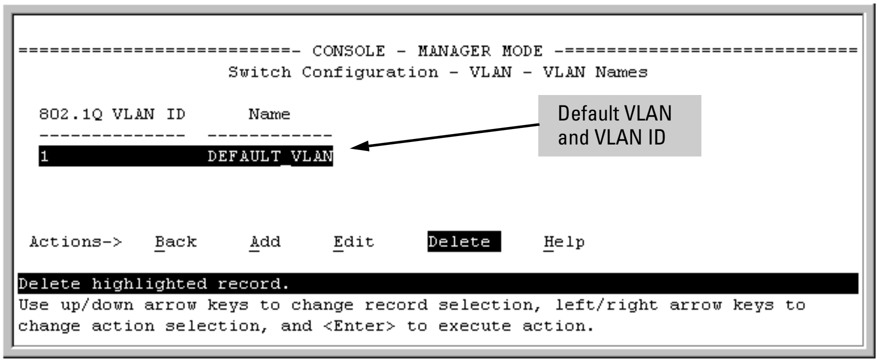

```
=========================- CONSOLE - MANAGER MODE -===========================
                  Switch Configuration - VLAN - VLAN Names

  802.1Q VLAN ID      Name
  --------------   ------------
  1                DEFAULT_VLAN




 Actions->   Back     Add      Edit      Delete     Help

Delete highlighted record.
Use up/down arrow keys to change record selection, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Default VLAN and VLAN ID

Figure 11-8. The Default VLAN Names Screen

2. Press A (for **Add**). You will then be prompted for a new VLAN name and VLAN ID:

   802.1Q VLAN ID : 1
   Name : _

3. Type in a VID (VLAN ID number). This can be any number from 2 to 4094 that is not already being used by another VLAN. (The switch reserves "1" for the default VLAN.)

   Remember that a VLAN *must* have the same VID in every switch in which you configure that same VLAN. (GVRP dynamically extends VLANs with correct VID numbering to other switches. See "GVRP" on page 11-34.)

4. Press ↓ to move the cursor to the **Name** line and type the VLAN name (up to 12 characters, with no spaces) of a new VLAN that you want to add, then press Enter.
   (Avoid these characters in VLAN names: **2**, **#**, **$**, **^**, **&**, *****, **(**, and **)**.)

5. Press S (for **Save**). You will then see the VLAN Names screen with the new VLAN listed.

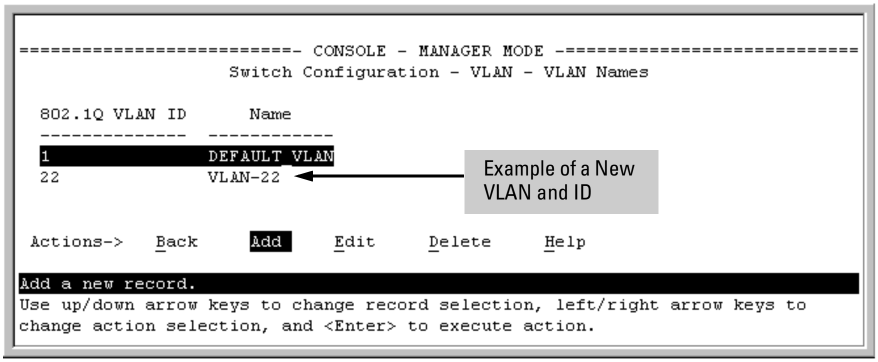

```
=========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - VLAN - VLAN Names

  802.1Q VLAN ID      Name
  --------------  ------------
  1               DEFAULT_VLAN
  22              VLAN-22

 Actions->  Back      Add     Edit      Delete      Help

Add a new record.
Use up/down arrow keys to change record selection, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Example of a New VLAN and ID

Figure 11-9. Example of VLAN Names Screen with a New VLAN Added

6. Repeat steps 2 through 5 to add more VLANs.

   Remember that you can add VLANs until you reach the number specified in the **Maximum VLANs to support** field on the VLAN Support screen (see figure 11-6 on page 11-10). This includes any VLANs added dynamically due to GVRP operation.

7. Return to the VLAN Menu to assign ports to the new VLAN(s) as described in the next section, "Adding or Changing a VLAN Port Assignment".

Adding or Changing a VLAN Port Assignment

Use this procedure to add ports to a VLAN or to change the VLAN assignment(s) for any port. (Ports not specifically assigned to a VLAN are automatically in the default VLAN.)

1. From the Main Menu select:

   2. Switch Configuration

   8. VLAN Menu . . .

   3. VLAN Port Assignment

   You will then see a VLAN Port Assignment screen similar to the following:

Default: In this example, the "VLAN-22" has been defined, but no ports have yet been assigned to it. ("No" means the port is not assigned to that VLAN.)
Using GVRP? If you plan on using GVRP, any ports you don't want to join should be changed to "Forbid".

A port can be assigned to several VLANs, but only one of those assignments can be "Untagged".

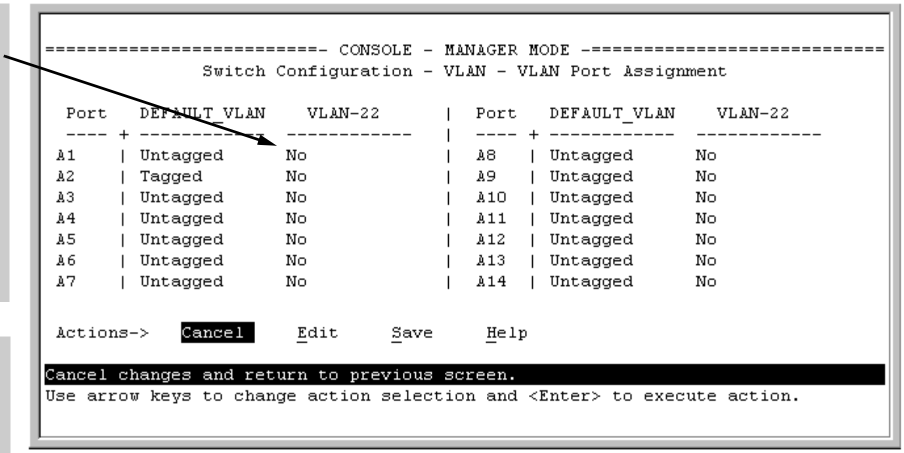

```
=========================- CONSOLE - MANAGER MODE -===========================
              Switch Configuration - VLAN - VLAN Port Assignment

  Port   DEFAULT_VLAN    VLAN-22       |  Port   DEFAULT_VLAN    VLAN-22
  ---- + ------------  ------------    |  ---- + ------------  ------------
 A1    | Untagged      No              |  A8   | Untagged      No
 A2    | Tagged        No              |  A9   | Untagged      No
 A3    | Untagged      No              |  A10  | Untagged      No
 A4    | Untagged      No              |  A11  | Untagged      No
 A5    | Untagged      No              |  A12  | Untagged      No
 A6    | Untagged      No              |  A13  | Untagged      No
 A7    | Untagged      No              |  A14  | Untagged      No

 Actions->  Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 11-10. Example of VLAN Port Assignment Screen

2. To change a port's VLAN assignment(s):
   a. Press [E] (for **Edit**).
   b. Use the arrow keys to select a VLAN assignment you want to change.
   c. Press the Space bar to make your assignment selection (**No**, **Tagged**, **Untagged**, or **Forbid**).

Note

For GVRP Operation: If you enable GVRP on the switch, "**No**" converts to "**Auto**", which allows the VLAN to dynamically join an advertised VLAN that has the same VID. See "Per-Port Options for Dynamic VLAN Advertising and Joining" on page 11-39.

Untagged VLANs: Only one untagged VLAN is allowed per port. Also, there must be at least one VLAN assigned to each port. In the factory default configuration, all ports are assigned to the default VLAN (DEFAULT_VLAN).

For example, if you want ports A4 and A5 to belong to both DEFAULT_VLAN and VLAN-22, and ports A6 and A7 to belong only to VLAN-22, you would use the settings in figure page 11-15. (This example assumes the default GVRP setting—disabled—and that you do not plan to enable GVRP later.)

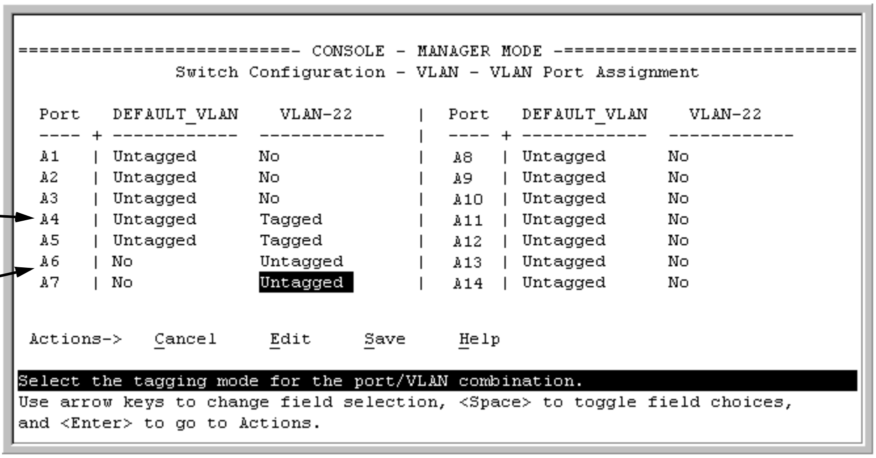

```
==========================- CONSOLE - MANAGER MODE -============================
               Switch Configuration - VLAN - VLAN Port Assignment

  Port    DEFAULT_VLAN    VLAN-22      |  Port    DEFAULT_VLAN    VLAN-22
  ---- + ------------  ------------    |  ---- + ------------  ------------
  A1   | Untagged      No              |  A8   | Untagged      No
  A2   | Untagged      No              |  A9   | Untagged      No
  A3   | Untagged      No              |  A10  | Untagged      No
  A4   | Untagged      Tagged          |  A11  | Untagged      No
  A5   | Untagged      Tagged          |  A12  | Untagged      No
  A6   | No            Untagged        |  A13  | Untagged      No
  A7   | No            Untagged        |  A14  | Untagged      No


 Actions->   Cancel     Edit     Save     Help

Select the tagging mode for the port/VLAN combination.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Ports A4 and A5 are assigned to both VLANs.

Ports A6 and A7 are assigned only to VLAN-22.

All other ports are assigned only to the Default VLAN.

Figure 11-11. Example of VLAN Assignments for Specific Ports

For information on VLAN tags ("Untagged" and "Tagged"), refer to "VLAN Tagging Information" on page 11-22.

d. If you are finished assigning ports to VLANs, press [Enter] and then [S] (for **Save**) to activate the changes you've made and to return to the Configuration menu. (The console then returns to the VLAN menu.)

3. Return to the Main menu.

CLI: Configuring VLAN Parameters

In the factory default state, all ports on the switch belong to the default VLAN (DEFAULT_VLAN) and are in the same broadcast/multicast domain. (The default VLAN is also the default primary VLAN—see "The Primary VLAN" on page 11-6.) You can configure up to 29 additional static VLANs by adding new VLAN names, and then assigning one or more ports to each VLAN. (The switch accepts a maximum of 256 VLANs, including the default VLAN and any dynamic VLANs the switch creates if you enable GVRP—page 11-34.) Note that each port can be assigned to multiple VLANs by using VLAN tagging. (See "VLAN Tagging Information" on page 11-22.)

VLAN Commands	**Page**
show vlans	below
show vlan <*vlan-id*>	11-17
max-vlans <1..256>	11-18
primary-vlan <*vlan-id*>	11-18
[no] vlan <*vlan-id*>	11-19
name <*vlan-name*>	11-20
[no] tagged <*port-list*>	11-20
[no] untagged <*port-list*>	11-20
[no] forbid	11-20
auto <*port-list*>	11-20 (Available if GVRP enabled.)
static-vlan <*vlan-id*>	11-20 (Available if GVRP enabled.)

Displaying the Switch's VLAN Configuration. The next command lists the VLANs currently running in the switch, with VID, VLAN name, and VLAN status. Dynamic VLANs appear only if the switch is running with GVRP enabled and one or more ports has dynamically joined an advertised VLAN. (In the default configuration, GVRP is disabled. (See "GVRP" on page 11-34.)

Syntax: show vlan

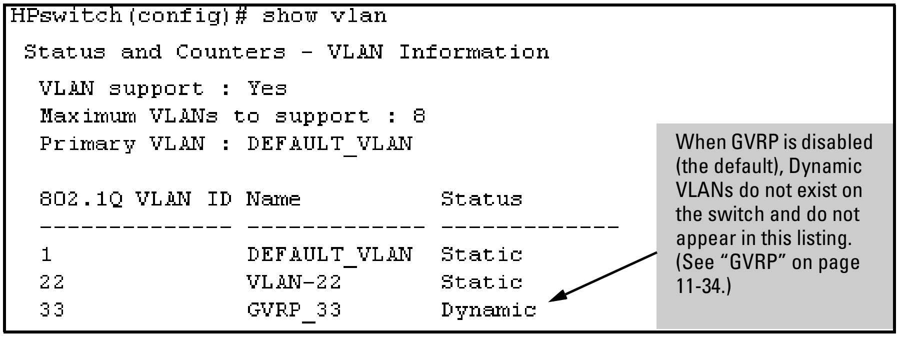

```
HPswitch(config)# show vlan

 Status and Counters - VLAN Information

  VLAN support : Yes
  Maximum VLANs to support : 8
  Primary VLAN : DEFAULT_VLAN

  802.1Q VLAN ID Name          Status
  -------------- ------------- -------------
  1              DEFAULT_VLAN  Static
  22             VLAN-22       Static
  33             GVRP_33       Dynamic
```

When GVRP is disabled (the default), Dynamic VLANs do not exist on the switch and do not appear in this listing. (See "GVRP" on page 11-34.)

Figure 11-12. Example of "Show VLAN" Listing (GVRP Enabled)

Displaying the Configuration for a Particular VLAN . This command uses the VID to identify and display the data for a specific static or dynamic VLAN.

Syntax: show vlan <*vlan-id*>

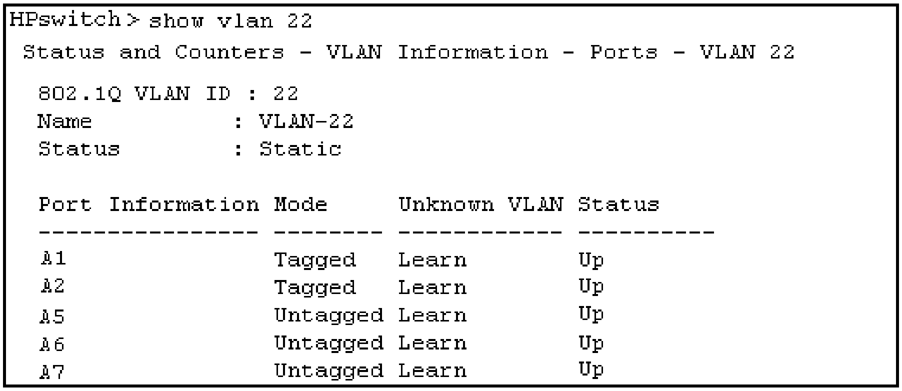

```
HPswitch> show vlan 22
 Status and Counters - VLAN Information - Ports - VLAN 22

  802.1Q VLAN ID : 22
  Name           : VLAN-22
  Status         : Static

  Port Information Mode     Unknown VLAN Status
  ---------------- -------- ------------ ----------
  A1               Tagged   Learn        Up
  A2               Tagged   Learn        Up
  A5               Untagged Learn        Up
  A6               Untagged Learn        Up
  A7               Untagged Learn        Up
```

Figure 11-13. Example of “Show VLAN” for a Specific Static VLAN

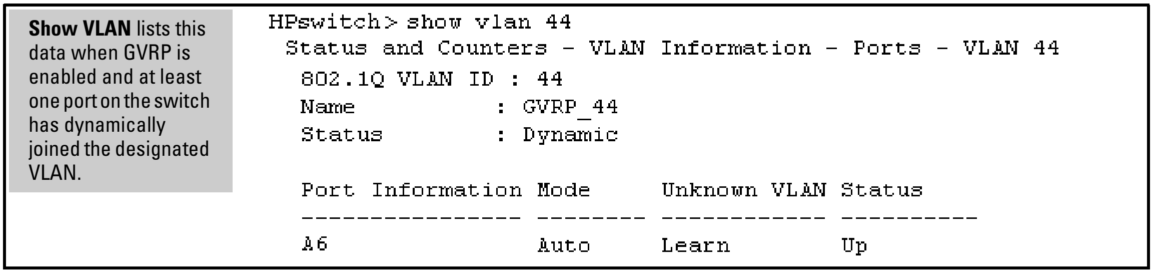

Show VLAN lists this data when GVRP is enabled and at least one port on the switch has dynamically joined the designated VLAN.

```
HPswitch> show vlan 44
 Status and Counters - VLAN Information - Ports - VLAN 44
  802.1Q VLAN ID : 44
  Name           : GVRP_44
  Status         : Dynamic

  Port Information Mode     Unknown VLAN Status
  ---------------- -------- ------------ ----------
  A6               Auto     Learn        Up
```

Figure 11-14. Example of “Show VLAN” for a Specific Dynamic VLAN

Changing the Number of VLANs Allowed on the Switch. By default, the switch allows a maximum of 8 VLANs. You can specify any value from 1 to 256. (If GVRP is enabled, this setting includes any dynamic VLANs on the switch.) As part of implementing a new value, you must execute a write memory command (to save the new value to the startup-config file) and then reboot the switch.

Syntax: max-vlans <1 .. 256>

For example, to reconfigure the switch to allow 10 VLANs:

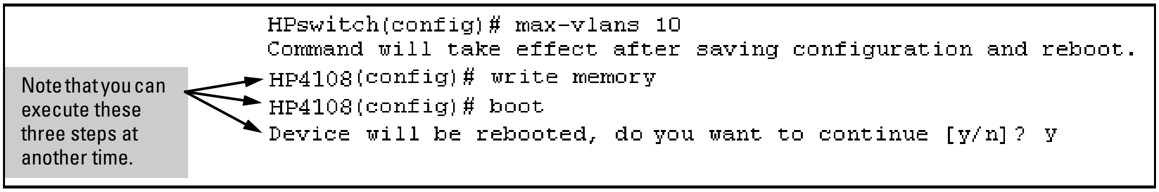

```
HPswitch(config)# max-vlans 10
Command will take effect after saving configuration and reboot.
HP4108(config)# write memory
HP4108(config)# boot
Device will be rebooted, do you want to continue [y/n]? y
```

Note that you can execute these three steps at another time.

Figure 11-15. Example of Command Sequence for Changing the Number of VLANs

Changing the Primary VLAN. In the factory-default configuration, the default VLAN (DEFAULT_VLAN) is the primary VLAN. However, you can designate any static VLAN on the switch as the primary VLAN. (For more on the primary VLAN, see "The Primary VLAN" on page 11-6.) To view the available VLANs and their respective VIDs, use **show vlan**.

Syntax: primary-vlan <*vlan-id*>

For example, to make VLAN 22 the primary VLAN:

```
HPswitch(config)# primary-vlan 22
```

Creating a New Static VLAN
Changing the VLAN Context Level.

With this command, entering a new VID creates a new static VLAN. Entering the VID or name of an existing static VLAN places you in the context level for that VLAN.

Syntax: vlan <*vlan-id*> [name <*name-str*>]

Creates a new static VLAN if a VLAN with that VID does not already exist, and places you in that VLAN's context level. If you do not use the name option, the switch uses "VLAN" and the new VID to automatically name the VLAN.
If the VLAN already exists, the switch places you in the context level for that VLAN.

vlan <*vlan-name*>

Places you in the context level for that static VLAN.

For example, to create a new static VLAN with a VID of 100:

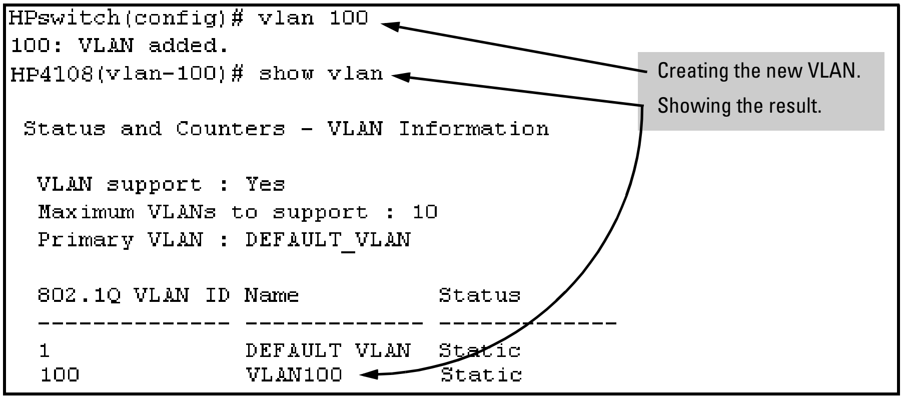

```
HPswitch(config)# vlan 100
100: VLAN added.
HP4108(vlan-100)# show vlan

 Status and Counters - VLAN Information

  VLAN support : Yes
  Maximum VLANs to support : 10
  Primary VLAN : DEFAULT_VLAN

  802.1Q VLAN ID Name          Status
  -------------- ------------- -------------
  1              DEFAULT VLAN  Static
  100            VLAN100       Static
```

Creating the new VLAN.

Showing the result.

Figure 11-16. Example of Creating a New Static VLAN

To go to a different VLAN context level, such as to the default VLAN:

```
HPswitch(vlan-100)# vlan default_vlank
HPswitch(vlan-1) _k
```

Converting a Dynamic VLAN to a Static VLAN. If GVRP is running on the switch and a port dynamically joins a VLAN, you can use the next command to convert the dynamic VLAN to a static VLAN. (For GVRP and dynamic VLAN operation, see "GVRP" on page 11-34.) This is necessary if you want to make the VLAN permanent. After you convert a dynamic VLAN to static, you must configure the switch's per-port participation in the VLAN in the same way that you would for any static VLAN.

Syntax: static-vlan <*vlan-id*> (Use **show vlan** to list current VIDs.)

For example, suppose a dynamic VLAN with a VID of 125 exists on the switch. The following command converts the VLAN to a static VLAN.

```
HPswitch(config)# static-vlan 125k
```

Configuring Static VLAN Name and Per-Port Settings. The **vlan <*vlan-id*>** command, used with the options listed below, changes the name of an existing static VLAN and changes the per-port VLAN membership settings.

Note

You can use these options from the configuration level by beginning the command with **vlan <*vlan-id*>**, or from the context level of the specific VLAN.

Syntax:	name <*vlan-name*>	Changes the name of the existing static VLAN. (Avoid spaces and the following characters in the <*vlan-name*> entry: **2**, **#**, **$**, **^**, **&**, *****, **(**, and **)**.)
	[no] tagged <*port-list*>	Configures the indicated port(s) as **Tagged** for the specified VLAN. The "**no**" version sets the port(s) to either **No** or (if GVRP is enabled) to **Auto**.
	[no] untagged <*port-list*>	Configures the indicated port(s) as **Untagged** for the specified VLAN. The "**no**" version sets the port(s) to either **No** or (if GVRP is enabled) to **Auto**.
	[no] forbid <*port-list*>	Configures the indicated port(s) as "forbidden" to participate in the designated VLAN. The "**no**" version sets the port(s) to either **No** or (if GVRP is enabled) to **Auto**.
	auto <*port-list*>	Available if GVRP is enabled on the switch. Returns the per-port settings for the specified VLAN to **Auto** operation. Note that **Auto** is the default per-port setting for a static VLAN if GVRP is running on the switch.

(For information on dynamic VLAN and GVRP operation, see "GVRP" on page 11-34.)

For example, suppose you have a VLAN named VLAN100 with a VID of 100, and all ports are set to No for this VLAN. To change the VLAN name to "Blue_Team" and set ports 1-5 to Tagged, you could do so with these commands:

```
HPswitch(config)# vlan 100 name Blue_Teamk
HPswitch(config)# vlan 100 tagged 1-5k
```

To move to the vlan 100 context level and execute the same commands:

```
HPswitch(config)# vlan 100k
HPswitch(vlan-100)# name Blue_Teamk
HPswitch(vlan-100)# tagged 1-5k
```

Similarly, to change the tagged ports in the above examples to **No** (or **Auto**, if GVRP is enabled), you could use either of the following commands.

At the config level, use:

```
HPswitch(config)# no vlan 100 tagged 1-5k
```

- *or* -

At the VLAN 100 context level, use:

```
HPswitch(vlan-100)# no tagged 1-5k
```

Note

You cannot use these commands with dynamic VLANs. Attempting to do so results in the message "**VLAN already exists.**" and no change occurs.

Web: Viewing and Configuring VLAN Parameters

In the web browser interface you can do the following:

- Add VLANs
- Rename VLANs
- Remove VLANs
- Configure GVRP mode
- Select a new Primary VLAN

To configure static VLAN port parameters, you will need to use the menu interface (available by Telnet from the web browser interface) or the CLI.

1. Click on the Configuration tab.

Click on Vlan Configuration.

2. Click on Add/Remove VLANs.

For web-based Help on how to use the web browser interface screen, click on the ? button provided on the web browser screen.

VLAN Tagging Information

VLAN tagging enables traffic from more than one VLAN to use the same port. (Even when two or more VLANs use the same port they remain as separate domains and cannot receive traffic from each other without going through an external router.) As mentioned earlier, a "tag" is simply a unique VLAN identification number (VLAN ID, or VID) assigned to a VLAN at the time that you configure the VLAN name in the switch. In the Series 5300XL switches the tag can be any number from 1 to 4094 that is not already assigned to a VLAN. When you subsequently assign a port to a given VLAN, you must implement the VLAN tag (VID) if the port will carry traffic for more than one VLAN. Otherwise, the port VLAN assignment can remain "untagged" because the tag is not needed. On a given switch, this means you should use the "Untagged" designation for a port VLAN assignment where the port is connected to non 802.1Q-compliant device or is assigned to only one VLAN. Use the "Tagged" designation when the port is assigned to more than one VLAN or the port is connected to a device that *does* comply with the 802.1Q standard.

For example, if port A7 on an 802.1Q-compliant switch is assigned to only the Red VLAN, the assignment can remain "untagged" because the port will forward traffic only for the Red VLAN. However, if both the Red and Green VLANs are assigned to port A7, then at least one of those VLAN assignments must be "tagged" so that Red VLAN traffic can be distinguished from Green VLAN traffic. The following illustration shows this concept:

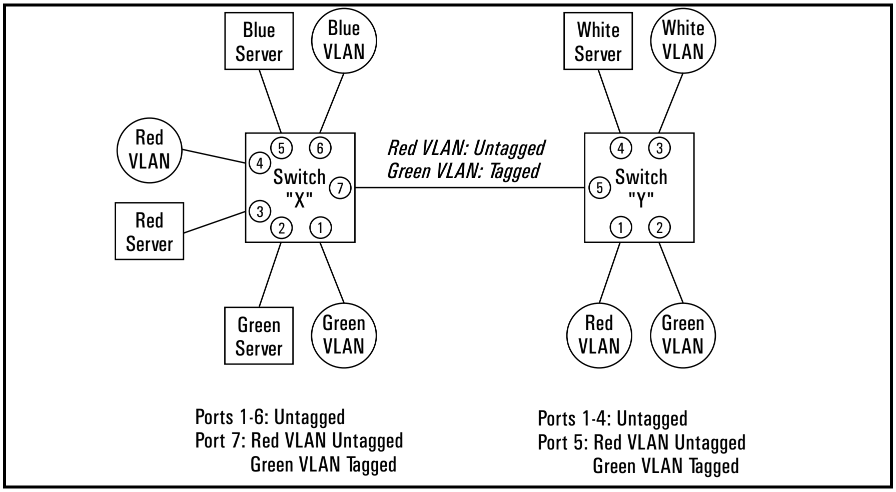

Figure 11-17. Example of Tagged and Untagged VLAN Port Assignments

- In switch X:
  - VLANs assigned to ports X1 - X6 can all be untagged because there is only one VLAN assignment per port. Red VLAN traffic will go out only the Red ports; Green VLAN traffic will go out only the Green ports, and so on. Devices connected to these ports do not have to be 802.1Q-compliant.
  - However, because both the Red VLAN and the Green VLAN are assigned to port X7, at least one of the VLANs must be tagged for this port.
- In switch Y:
  - VLANs assigned to ports Y1 - Y4 can all be untagged because there is only one VLAN assignment per port. Devices connected to these ports do not have to be 802.1Q-compliant.
  - Because both the Red VLAN and the Green VLAN are assigned to port Y5, at least one of the VLANs must be tagged for this port.
- In both switches: The ports on the link between the two switches must be configured the same. As shown in figure 11-17 (above), the Red VLAN must be untagged on port X7 and Y5 and the Green VLAN must be tagged on port X7 and Y5, or vice-versa.

Note

Each 802.1Q-compliant VLAN must have its own unique VID number, and that VLAN *must* be given the same VID in every device in which it is configured. That is, if the Red VLAN has a VID of 10 in switch X, then 10 must also be used for the Red VID in switch Y.

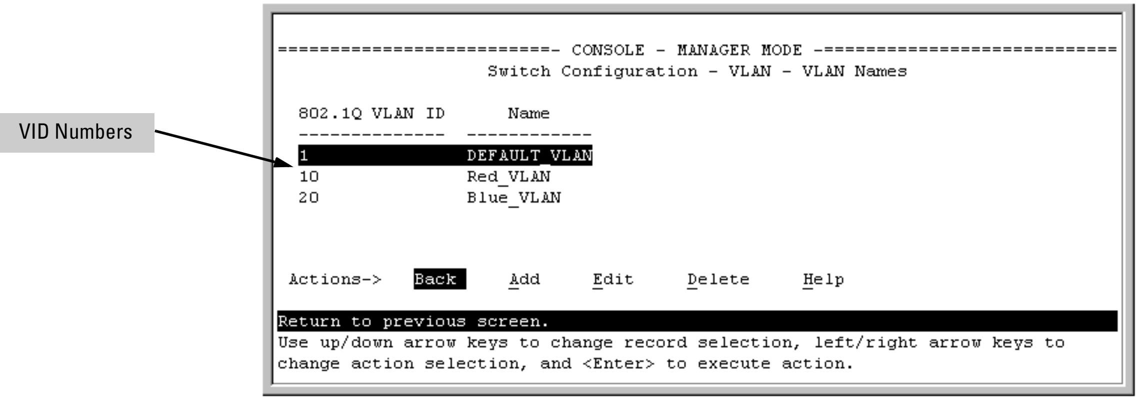

```
==========================- CONSOLE - MANAGER MODE -===========================
                      Switch Configuration - VLAN - VLAN Names

  802.1Q VLAN ID      Name
  --------------   ------------
  1                DEFAULT_VLAN
  10               Red_VLAN
  20               Blue_VLAN



 Actions->   Back     Add     Edit     Delete     Help

Return to previous screen.
Use up/down arrow keys to change record selection, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure 11-18. Example of VLAN ID Numbers Assigned in the VLAN Names Screen

VLAN tagging gives you several options:

- Since the purpose of VLAN tagging is to allow multiple VLANs on the same port, any port that has only one VLAN assigned to it can be configured as "Untagged" (the default).
- Any port that has two or more VLANs assigned to it can have one VLAN assignment for that port as "Untagged". All other VLANs assigned to the same port must be configured as "Tagged". (There can be no more than one Untagged VLAN on a port.)
- If all end nodes on a port comply with the 802.1Q standard and are configured to use the correct VID, then, you can configure all VLAN assignments on a port as "Tagged" if doing so makes it easier to manage your VLAN assignments, or for security reasons.

For example, in the following network, switches X and Y and servers S1 and S2 are 802.1Q-compliant. (Server S3 could also be 802.1Q-compliant, but it makes no difference for this example.)

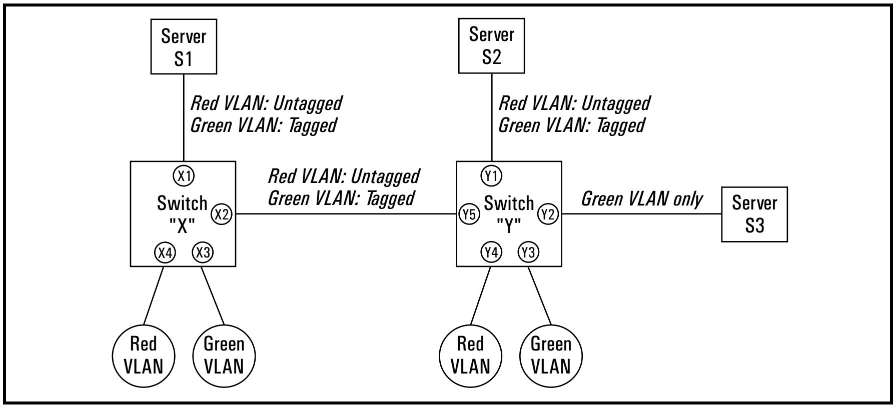

Figure 11-19. Example of Networked 802.1Q-Compliant Devices with Multiple VLANs on Some Ports

The VLANs assigned to ports X3, X4, Y2, Y3, and Y4 can all be untagged because there is only one VLAN assigned per port. Port X1 has multiple VLANs assigned, which means that one VLAN assigned to this port can be untagged and any others must be tagged. The same applies to ports X2, Y1, and Y5.

Switch X				Switch Y		
Port	Red VLAN	Green VLAN		Port	Red VLAN	Green VLAN
X1	Untagged	Tagged		Y1	Untagged	Tagged
X2	Untagged	Tagged		Y2	No*	Untagged
X3	No*	Untagged		Y3	No*	Untagged
X4	Untagged	No*		Y4	Untagged	No*
				Y5	Untagged	Tagged

*"No" means the port is not a member of that VLAN. For example, port X3 is not a member of the Red VLAN and does not carry Red VLAN traffic. Also, if GVRP were enabled, "Auto" would appear instead of "No".

Note

VLAN configurations on ports connected by the same link must match. Because ports X2 and Y5 are opposite ends of the same point-to-point connection, both ports must have the same VLAN configuration; that is, both ports configure the Red VLAN as "Untagged" and the Green VLAN as "Tagged".

To summarize:

VLANs Per Port	Tagging Scheme
1	Untagged or Tagged. If the device connected to the port is 802.1Q-compliant, then the recommended choice is "Tagged".
2 or More	1 VLAN Untagged; all others Tagged or All VLANs Tagged
A given VLAN *must* have the same VID on any 802.1Q-compliant device in which the VLAN is configured.	
The ports connecting two 802.1Q devices should have identical VLAN configurations, as shown for ports X2 and Y5, above.	

The Secure Management VLAN

Configures a secure Management VLAN by creating an isolated network for managing the HP Procurve switches that support this feature. (As of June 1, 2002, includes the HP Procurve Series 5300XL switches and Series 5300XL switches.) Access to this VLAN, and to the switch's management functions (Menu, CLI, and web browser interface) is available only through ports configured as members.

- Multiple ports on the switch can belong to the Management VLAN. This allows connections for multiple management stations you want to have access to the Management VLAN, while at the same time allowing Management VLAN links between switches configured for the same Management VLAN.
- Only traffic from the Management VLAN can manage the switch, which means that only the workstations and PCs connected to ports belonging to the Management VLAN can manage and reconfigure the switch.

Figure 11-20 illustrates use of the Management VLAN feature to support management access by a group of management workstations.

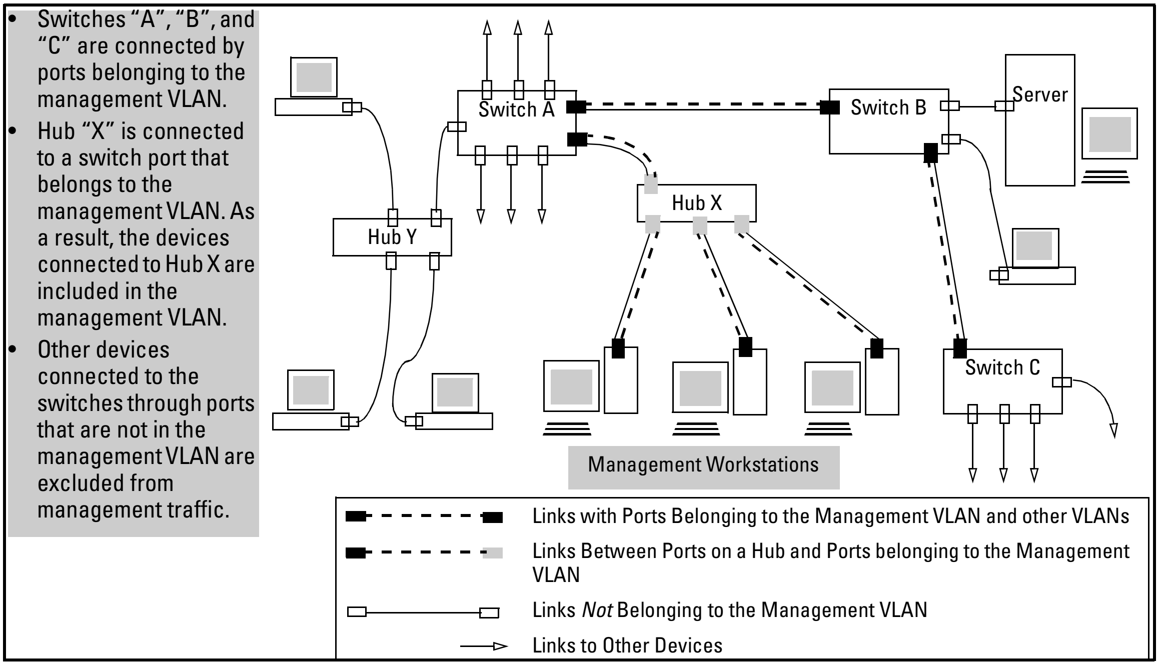

- Switches "A", "B", and "C" are connected by ports belonging to the management VLAN.
- Hub "X" is connected to a switch port that belongs to the management VLAN. As a result, the devices connected to Hub X are included in the management VLAN.
- Other devices connected to the switches through ports that are not in the management VLAN are excluded from management traffic.

Switch A
Switch B
Server
Hub X
Hub Y
Switch C
Management Workstations
Links with Ports Belonging to the Management VLAN and other VLANs
Links Between Ports on a Hub and Ports belonging to the Management VLAN
Links *Not* Belonging to the Management VLAN
Links to Other Devices

Figure 11-20. Example of Potential Security Breaches

In figure 11-21, Workstation 1 has management access to all three switches through the Management VLAN, while the PCs do not. This is because configuring a switch to recognize a Management VLAN automatically excludes attempts to send management traffic from any other VLAN.

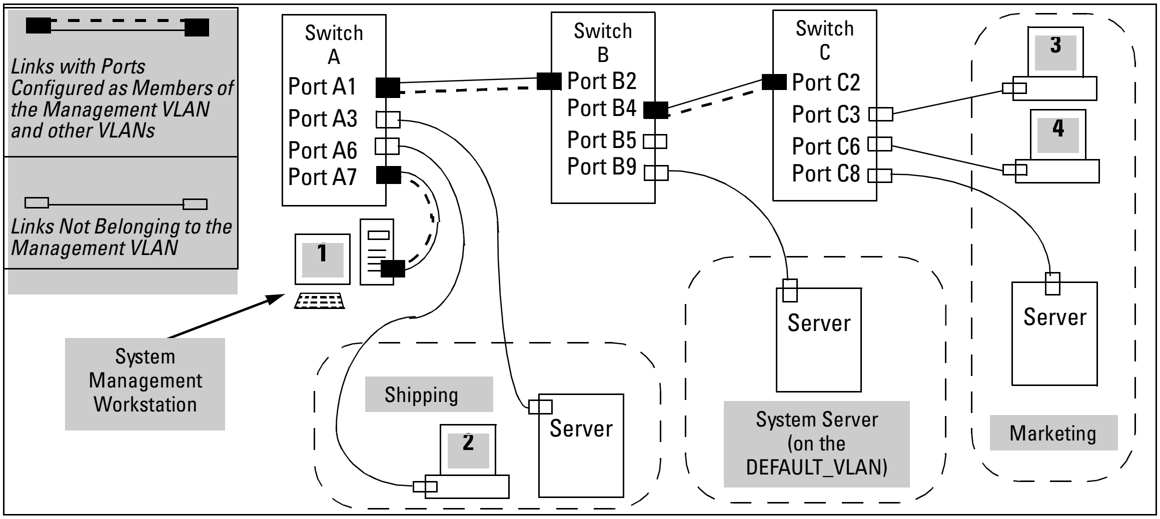
Links with Ports Configured as Members of the Management VLAN and other VLANs
Links Not Belonging to the Management VLAN
System Management Workstation
Switch A
Port A1
Port A3
Port A6
Port A7
1
Switch B
Port B2
Port B4
Port B5
Port B9
Switch C
Port C2
Port C3
Port C6
Port C8
3
4
Shipping
2
Server
Server
System Server (on the DEFAULT_VLAN)
Server
Marketing

Figure 11-21. Example of Management VLAN Control in a LAN

Table 11-2.VLAN Membership in Figure 11-21

Switch	A1	A3	A6	A7	B2	B4	B5	B9	C2	C3	C6	C8
Management VLAN (VID = 7)	**Y**	N	N	**Y**	**Y**	**Y**	N	N	**Y**	N	N	N
Marketing VLAN (VID = 12)	N	N	N	N	N	N	N	N	N	**Y**	**Y**	**Y**
Shipping Dept. VLAN (VID = 20)	N	**Y**	**Y**	N	N	N	N	N	N	N	N	N
DEFAULT-VLAN (VID = 1)	**Y**	**Y**	**Y**	**Y**	**Y**	**Y**	**Y**	**Y**	**Y**	**Y**	**Y**	**Y**

Preparation

1. Determine a VID and VLAN name suitable for your Management VLAN.
2. Determine the IP addressing for the Management VLAN (**DHCP/Bootp** or **Manual**.
3. Plan your Management VLAN topology to use HP Procurve switches that support this feature. (As of June 1, 2002, this includes the HP Procurve Series 5300XL and Series 4100GL switches.) The ports belonging to the Management VLAN should be only the following:
   - Ports to which you will connect authorized management stations (such as Port A7 in figure 11-21.)

- Ports on one switch that you will use to extend the Management VLAN to ports on other HP Procurve switches (such as ports A1 and B2 or B4 and C2 in figure 11-21 on page 11-28.).

Hubs dedicated to connecting management stations to the Management VLAN can also be included in the above topology. Note that any device connected to a hub in the Management VLAN will also have Management VLAN access.

4. Configure the Management VLAN on the selected switch ports.
5. Test the management VLAN from all of the management stations authorized to use the Management VLAN, including any SNMP-based network management stations. Ensure that you include testing any Management VLAN links between switches.

Note

If you configure a Management VLAN on a switch by using a Telnet connection through a port that is not in the Management VLAN, then you will lose management contact with the switch if you log off your Telnet connection or execute **write memory** and **reboot** the switch.

Configuration

Syntax: [no] management-vlan < *vlan-id* | *vlan-name* >
show vlan-info

Default: Disabled

For example, suppose you have already configured a VLAN named **My_VLAN** with a VID of 100 in a Series 5300XL switch. Now you want to configure the switch to do the following:

- Use **My_VLAN** as a Management VLAN (tagged, in this case) to connect port A1 on switch "A" to a management station. (The management station includes a network interface card with 802.1Q tagged VLAN capability.)
- Use port A2 to extend the Management VLAN to port B1 (which is already configured as a tagged member of **My_VLAN**) on an adjacent Series 5300XL switch.

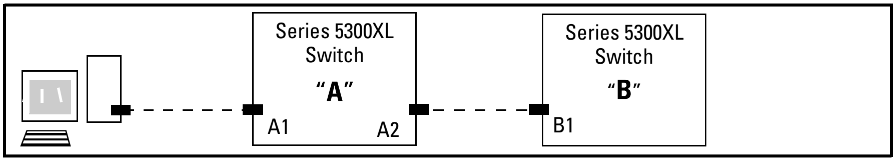

Figure 11-22. Illustration of Configuration Example

```
HPswitch (config)# management-vlan 100k
HPswitch (config)# vlan 100 tagged a1k
HPswitch (config)# vlan 100 tagged a2
```

Deleting the Management VLAN. You can disable the Secure Management feature without deleting the VLAN itself. For example, either of the following commands disables the Secure Management feature in the above example:

```
HPswitch (config)# no management-vlan 100k
HPswitch (config)# no management-vlan my_vlank
```

Operating Notes for Management VLANs

- On Series 5300XL switches with routing enabled, routing between the Management VLAN and other VLANs is not allowed.
- If there are more than 25 VLANs configured on the switch, reboot the switch after configuring the management VLAN. (*HP Series 5300XL switches only.*)
- If you implement a Management VLAN in a switch mesh environment, all meshed ports on Series 5300XL switches will be members of the Management VLAN.
- Only one Management-VLAN can be active in the switch. If one Management-VLAN VID is saved in the startup-config file and you configure a different VID in the running-config file, the switch uses the running-config version until you either use the **write-memory** command or reboot the switch.
- During a Telnet session to the switch, if you configure the Management-VLAN to a VID that excludes the port through which you are connected to the switch, you will continue to have access only until you terminate the session by logging out or rebooting the switch.
- During a web browser session to the switch, if you configure the Management-VLAN to a VID that excludes the port through which you are connected to the switch, you will continue to have access only until you close the browser session or rebooting the switch.

Note

The Management-VLAN feature does not control management access through a direct connection to the switch's serial port.

- Enabling Spanning Tree where there are multiple links using separate VLANs, including the Management VLAN, between a pair of switches, Spanning Tree will force the blocking of one or more links. This may include the link carrying the Management VLAN, which will cause loss of management access to some devices. This can also occur where meshing is configured and the Management VLAN is configured on a separate link.

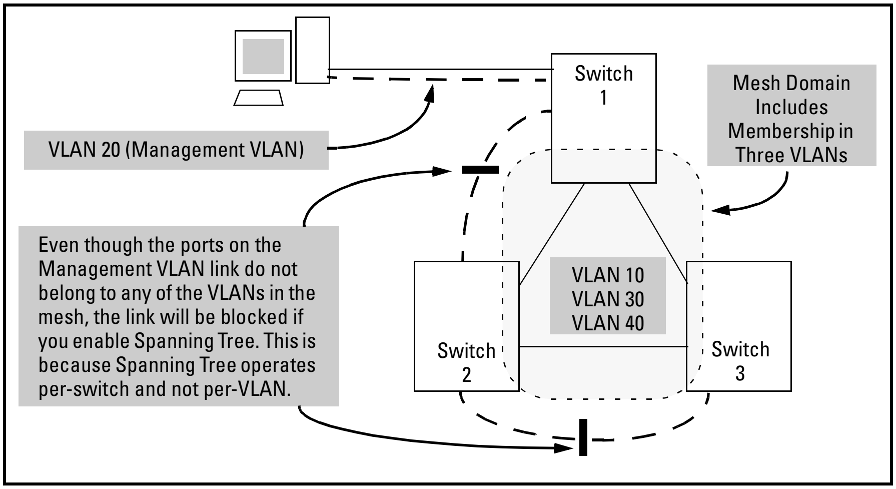
Switch 1
Mesh Domain Includes Membership in Three VLANs
VLAN 20 (Management VLAN)
Even though the ports on the Management VLAN link do not belong to any of the VLANs in the mesh, the link will be blocked if you enable Spanning Tree. This is because Spanning Tree operates per-switch and not per-VLAN.
VLAN 10
VLAN 30
VLAN 40
Switch 2
Switch 3

Figure 11-23. Example of Inadvertently Blocking a Management VLAN Link by Implementing Spanning Tree

Effect of VLANs on Other Switch Features

Spanning Tree Operation with VLANs

Because the Series 5300XL switches follows the 802.1Q VLAN recommendation to use single-instance spanning tree, Spanning Tree operates across all ports on the switch (regardless of VLAN assignments) instead of on a per-VLAN basis. This means that if redundant physical links exist between the switch and another 802.1Q device, all but one link will be blocked, regardless of whether the redundant links are in separate VLANs. However, you can use port trunking to prevent Spanning Tree from unnecessarily blocking ports (and to improve overall network performance). Refer to “STP Operation with 802.1Q VLANs” on page 13-4.

Note that Spanning Tree operates differently in different devices. For example, in the (obsolete, non-802.1Q) HP Switch 2000 and the HP Switch 800T, Spanning Tree operates on a per-VLAN basis, allowing redundant physical links as long as they are in separate VLANs.

IP Interfaces

There is a one-to-one relationship between a VLAN and an IP network interface. Since the VLAN is defined by a group of ports, the state (up/down) of those ports determines the state of the IP network interface associated with that VLAN. When a VLAN comes up because one or more of its ports is up, the IP interface for that VLAN is also activated. Likewise, when a VLAN is deactivated because all of its ports are down, the corresponding IP interface is also deactivated.

VLAN MAC Address

The Series 5300XL switches have one unique MAC address for all of their VLAN interfaces. You can send an 802.2 test packet to this MAC address to verify connectivity to the switch. Likewise, you can assign an IP address to the VLAN interface, and when you Ping that address, ARP will resolve the IP address to this single MAC address.

Port Trunks

When assigning a port trunk to a VLAN, all ports in the trunk are automatically assigned to the same VLAN. You cannot split trunk members across multiple VLANs. Also, a port trunk is tagged, untagged, or excluded from a VLAN in the same way as for individual, untrunked ports.

Port Monitoring

If you designate a port on the switch for network monitoring, this port will appear in the Port VLAN Assignment screen and can be configured as a member of any VLAN. For information on how broadcast, multicast, and unicast packets are tagged inside and outside of the VLAN to which the monitor port is assigned, see “VLAN-Related Problems” on page C-19.

VLAN Restrictions

- A port must be a member of at least one VLAN. In the factory default configuration, all ports are assigned to the default VLAN (DEFAULT_VLAN; VID = 1).
- A port can be assigned to several VLANs, but only one of those assignments can be untagged. (The “Untagged” designation enables VLAN operation with non 802.1Q-compliant devices.)
- An external router must be used to communicate between tagged VLANs on the switch.

- Before you can delete a VLAN, you must first re-assign all ports in the VLAN to another VLAN.

HP Router Requirements. *Use the Hewlett-Packard version A.09.70 (or later) router OS release if any of the following Hewlett-Packard routers are installed in networks in which you will be using VLANs:*

HP Router 440 (formerly Router ER)
HP Router 470 (formerly Router LR)
HP Router 480 (formerly Router BR)
HP Router 650

Release A.09.74 is available on the World Wide Web at

http://www.hp.com/go/hpprocurve

Click on **software**, then **routers**.

GVRP

Feature	Default	Menu	CLI	Web
view GVRP configuration	n/a	page 11-42	page 11-44	page 11-47
list static and dynamic VLANs on a GVRP-enabled switch	n/a	—	page 11-46	page 11-47
enable or disable GVRP	disabled	page 11-42	page 11-45	page 11-47
enable or disable GVRP on individual ports	enabled	page 11-42	page 11-45	—
control how individual ports will handle advertisements for new VLANs	Learn	page 11-42	page 11-45	page 11-47
convert a dynamic VLAN to a static VLAN	n/a	—	page 11-47	—
configure static VLANs	DEFAULT_VLAN (VID = 1)	page 11-10	page 11-15	page 11-47

GVRP—GARP VLAN Registration Protocol—is an application of the Generic Attribute Registration Protocol—GARP. GVRP is defined in the IEEE 802.1Q standard, and GARP is defined in the IEEE 802.1D-1998 standard.

Note

To understand and use GVRP you must have a working knowledge of 802.1Q VLAN tagging. (See "Port-Based Virtual LANs (Static VLANs)" on page 11-3.)

GVRP uses "GVRP Bridge Protocol Data Units" ("GVRP BPDUs") to "advertise" static VLANs. In this manual, a GVRP BPDU is termed an *advertisement*. Advertisements are sent outbound from ports on a switch to the devices directly connected to those ports.

GVRP enables the Series 5300XL switches to dynamically create 802.1Q-compliant VLANs on links with other devices running GVRP. This enables the switch to automatically create VLAN links between GVRP-aware devices. (A GVRP link can include intermediate devices that are not GVRP-aware.) This operation reduces the chances for errors in VLAN configuration by automatically providing VLAN ID (VID) consistency across the network. That is, you can use GVRP to propagate VLANs to other GVRP-aware devices instead of manually having to set up VLANs across your network. After the switch creates a dynamic VLAN, you can optionally use the CLI **static <*vlan-id*>** command to convert it to a static VLAN or allow it to continue as a dynamic VLAN for as long as needed. You can also use GVRP to dynamically enable port membership in static VLANs configured on a switch.

General Operation

When GVRP is enabled on a switch, the VID for any static VLANs configured on the switch is *advertised* (using BPDUs—Bridge Protocol Data Units) out all ports, regardless of whether a port is up or assigned to any particular VLAN. A GVRP-aware port on another device that receives the advertisements over a link can dynamically join the advertised VLAN.

A dynamic VLAN (that is, a VLAN learned through GVRP) is tagged on the port on which it was learned. Also, a GVRP-enabled port can forward an advertisement for a VLAN it learned about from other ports on the same switch (internal source), but the forwarding port will not itself join that VLAN until an advertisement for that VLAN is received through a link from another device (external source) on that specific port

Operating Note: When a GVRP-aware port on a switch learns a VID through GVRP from another device, the switch begins advertising that VID out all of its ports except the port on which the VID was learned.

Core switch with static VLANs (VID= 1, 2, & 3). Port 2 is a member of VIDs 1, 2, & 3.

1. Port 2 advertises VIDs 1, 2, & 3.

2. Port 1 receives advertisement of VIDs 1, 2, & 3 AND becomes a member of VIDs 1, 2, & 3.

3. Port 3 advertises VIDs 1, 2, & 3, but port 3 is NOT a member of VIDs 1, 2, & 3 at this point.

4. Port 4 receives advertisement of VIDs 1, 2, & 3 AND becomes a member of VIDs 1, 2, & 3.

5. Port 5 advertises VIDs 1, 2, & 3, but port 5 is NOT a member of VIDs 1, 2, & 3 at this point.

Port 6 is statically configured to be a member of VID 3.

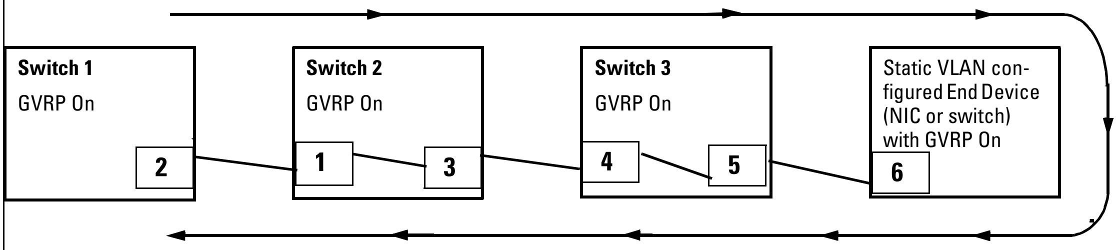
Switch 1
GVRP On
2
Switch 2
GVRP On
1
3
Switch 3
GVRP On
4
5
Static VLAN configured End Device (NIC or switch) with GVRP On
6

11. Port 2 receives advertisement of VID 3. (Port 2 is already statically configured for VID 3.)

9. Port 3 receives advertisement of VID 3 AND becomes a member of VID 3. (Still not a member of VIDs 1 & 2.)

10. Port 1 advertises VID 3.

7. Port 5 receives advertisement of VID 3 AND becomes a member of VID 3. (Still not a member of VIDs 1 & 2.)

8. Port 4 advertises VID 3.

6. Port 6 advertises VID 3.

Figure 11-24. .Example of Forwarding Advertisements and Dynamic Joining

Note that if a static VLAN is configured on at least one port of a switch, and that port has established a link with another device, then all other ports of that switch will send advertisements for that VLAN.

For example, in the following figure, Tagged VLAN ports on switch "A" and switch "C" advertise VLANs 22 and 33 to ports on other GVRP-enabled switches that can dynamically join the VLANs.

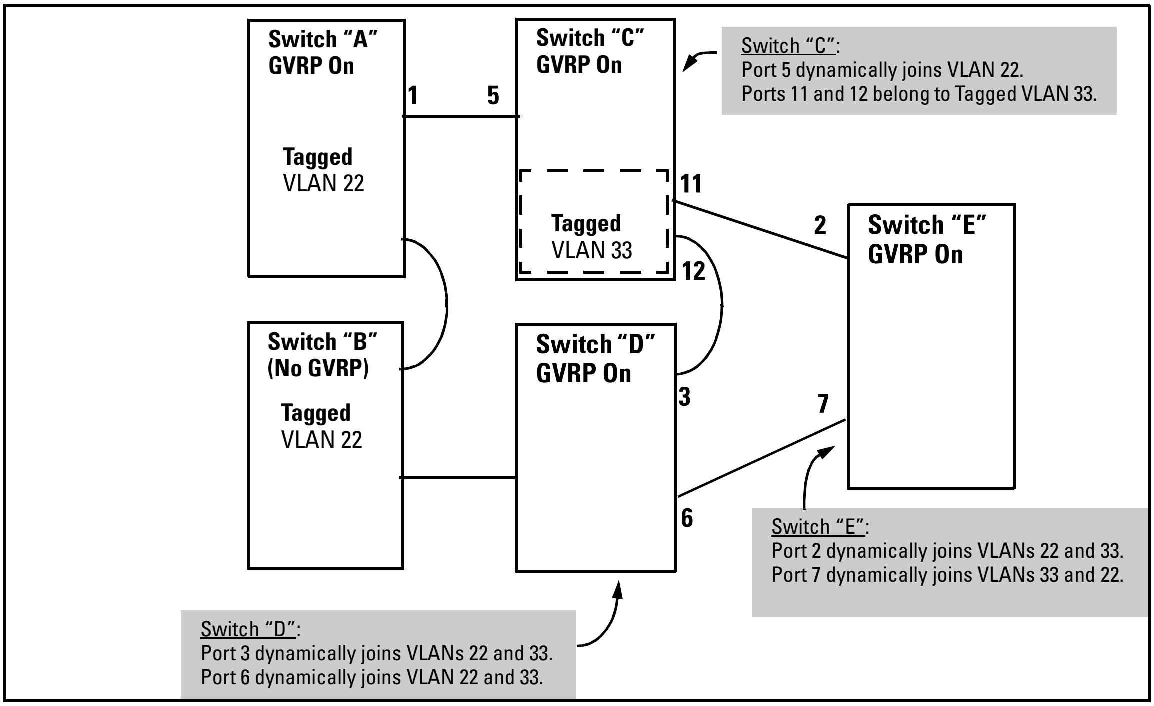

Switch "C":
Port 5 dynamically joins VLAN 22.
Ports 11 and 12 belong to Tagged VLAN 33.

Switch "E":
Port 2 dynamically joins VLANs 22 and 33.
Port 7 dynamically joins VLANs 33 and 22.

Switch "D":
Port 3 dynamically joins VLANs 22 and 33.
Port 6 dynamically joins VLAN 22 and 33.

Figure 11-25. Example of GVRP Operation

Note

A port can learn of a dynamic VLAN through devices that are not aware of GVRP (Switch "B", above). VLANs must be disabled in GVRP-unaware devices to allow tagged packets to pass through.

A GVRP-aware port receiving advertisements has these options:

- If there is not already a static VLAN with the advertised VID on the receiving port, then dynamically create the VLAN and become a member.
- If the switch already has a static VLAN assignment with the same VID as in the advertisement, and the port is configured to **Auto** for that VLAN, then the port will dynamically join the VLAN and begin moving that VLAN's traffic. (For more detail on **Auto**, see "Per-Port Options for Dynamic VLAN Advertising and Joining" on page 11-39.)
- Ignore the advertisement for that VID.
- Don't participate in that VLAN.

Note also that a port belonging to a Tagged or Untagged static VLAN has these configurable options:

- Send VLAN advertisements, and also receive advertisements for VLANs on other ports and dynamically join those VLANs.
- Send VLAN advertisements, but ignore advertisements received from other ports.
- Avoid GVRP participation by not sending advertisements and dropping any advertisements received from other devices.

IP Addressing. A dynamic VLAN does not have an IP address, and moves traffic on the basis of port membership in VLANs. However, after GVRP creates a dynamic VLAN, you can convert it to a static VLAN. Note that it is then necessary to assign ports to the VLAN in the same way that you would for a static VLAN that you created manually. In the static state you can configure IP addressing on the VLAN and access it in the same way that you would any other static (manually created) VLAN.

Per-Port Options for Handling GVRP "Unknown VLANs"

An "unknown VLAN" is a VLAN that the switch learns of by receiving an advertisement for that VLAN on a port that is not already a member of that VLAN. If the port is configured to learn unknown VLANs, then the VLAN is dynamically created and the port becomes a tagged member of the VLAN. For example, suppose that in figure 11-25 (page 11-36), port 1 on switch "A" is connected to port 5 on switch "C". Because switch "A" has VLAN 22 statically configured, while switch "C" does not have this VLAN statically configured (and does not "Forbid" VLAN 22 on port 5), VLAN 22 is handled as an "Unknown VLAN" on port 5 in switch "C". Conversely, if VLAN 22 was statically configured on switch C, but port 5 was not a member, port 5 would become a member when advertisements for VLAN 22 were received from switch "A".

When you enable GVRP on a switch, you have the per-port join-request options listed in table 11-3:

Table 11-3. Options for Handling "Unknown VLAN" Advertisements:

Unknown VLAN Mode	Operation
Learn (the Default)	Enables the port to become a member of any unknown VLAN for which it receives an advertisement. Allows the port to advertise other VLANs that have at least one other port on the same switch as a member.
Block	Prevents the port from joining any new dynamic VLANs for which it receives an advertisement. Allows the port to advertise other VLANs that have at least one other port as a member.
Disable	Causes the port to ignore and drop all GVRP advertisements it receives and also prevents the port from sending any GVRP advertisements.

The CLI **show gvrp** command and the menu interface VLAN Support screen show a switch's current GVRP configuration, including the Unknown VLAN settings.

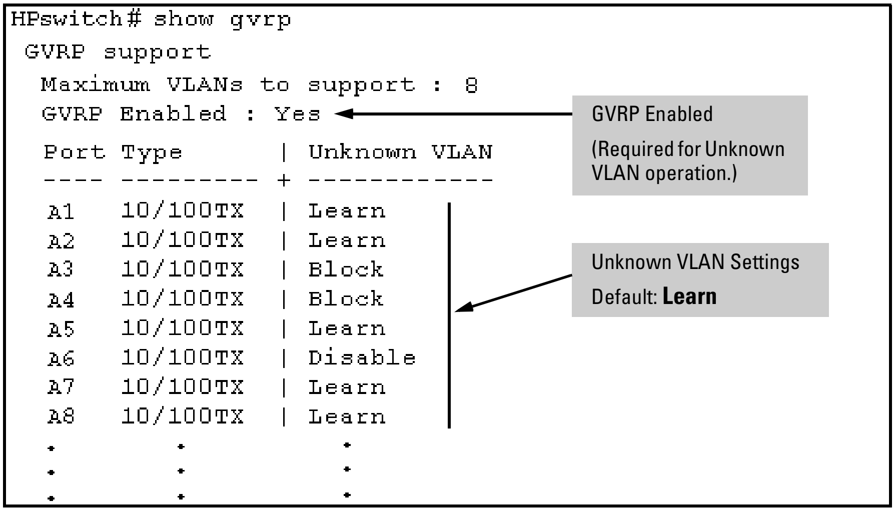

```
HPswitch# show gvrp
 GVRP support

  Maximum VLANs to support :  8
  GVRP Enabled : Yes

  Port Type       | Unknown VLAN
  ---- --------- + ------------
   A1   10/100TX  | Learn
   A2   10/100TX  | Learn
   A3   10/100TX  | Block
   A4   10/100TX  | Block
   A5   10/100TX  | Learn
   A6   10/100TX  | Disable
   A7   10/100TX  | Learn
   A8   10/100TX  | Learn
   .       .          .
   .       .          .
   .       .          .
```

GVRP Enabled

(Required for Unknown VLAN operation.)

Unknown VLAN Settings

Default: **Learn**

Figure 11-26. Example of GVRP Unknown VLAN Settings

Per-Port Options for Dynamic VLAN Advertising and Joining

Initiating Advertisements. As described in the preceding section, to enable dynamic joins, GVRP must be enabled and a port must be configured to Learn (the default). However, to send advertisements in your network, one or more static (**Tagged**, **Untagged**, or **Auto**) VLANs must be configured on one or more switches (with GVRP enabled), depending on your topology.

Enabling a Port for Dynamic Joins. You can configure a port to dynamically join a static VLAN. The join will then occur if that port subsequently receives an advertisement for the static VLAN. (This is done by using the **Auto** and **Learn** options described in table 11-4, below.

Parameters for Controlling VLAN Propagation Behavior. You can configure an individual port to actively or passively participate in dynamic VLAN propagation or to ignore dynamic VLAN (GVRP) operation. These options are controlled by the GVRP “Unknown VLAN” and the static VLAN configuration parameters, as described in the following table:

Table 11-4. Controlling VLAN Behavior on Ports with Static VLANs

Per-Port "Unknown VLAN" (GVRP) Configuration	Static VLAN Options—Per VLAN Specified on Each Port [1]		
	Port Activity: Tagged or Untagged (Per VLAN)[2]	**Port Activity: Auto[2] (Per VLAN)**	**Port Activity: Forbid (Per VLAN)[2]**
Learn (the Default)	The port: • Belongs to specified VLAN. • Advertises specified VLAN. • Can become a member of dynamic VLANs for which it receives advertisements. • Advertises dynamic VLANs that have at least one other port (on the same switch) as a member.	The port: • Will become a member of specified VLAN if it receives advertisements for specified VLAN from another device. • Will advertise specified VLAN. • Can become a member of other, dynamic VLANs for which it receives advertisements. • Will advertise a dynamic VLAN that has at least one other port (on the same switch) as a member.	The port: 1. Will not become a member of the specified VLAN. 1. Will not advertise specified VLAN. 1. Can become a member of other dynamic VLANs for which it receives advertisements. 1. Will advertise a dynamic VLAN that has at least one other port on the same switch as a member.
Block	The port: • Belongs to the specified VLAN. • Advertises this VLAN. • Will not become a member of new dynamic VLANs for which it receives advertisements. • Will advertise dynamic VLANs that have at least one other port as a member.	The port: • Will become a member of specified VLAN if it receives advertisements for this VLAN. • Will advertise this VLAN. • Will not become a member of new dynamic VLANs for which it receives advertisements. • Will advertise dynamic VLANs that have at least one other port (on the same switch) as a member.	The port: • Will not become a member of the specified VLAN. • Will not advertise this VLAN. • Will not become a member of dynamic VLANs for which it receives advertisements. • Will advertise dynamic VLANs that have at least one other port (on the same switch) as a member.
Disable	The port: • Is a member of the specified VLAN. • Will ignore GVRP PDUs. • Will not join any advertised VLANs. • Will not advertise VLANs.	The port: • Will not become a member of the specified VLAN. • Will ignore GVRP PDUs. • Will not join any dynamic VLANs. • Will not advertise VLANs.	The port: • Will not become a member of this VLAN. • Will ignore GVRP PDUs. • Will not join any dynamic VLANs. • Will not advertise VLANs.

[1] Each port of a Series 5300XL switches must be a Tagged or Untagged member of at least one VLAN. Thus, any port configured for GVRP to Learn or Block will generate and forward advertisements for static VLAN(s) configured on the switch and also for dynamic VLANs the switch learns on other ports..

[2] To configure tagging, **Auto**, or **Forbid**, see "Configuring Static VLAN Name and Per-Port Settings" on page 11-20 (for the CLI) or "Adding or Changing a VLAN Port Assignment" on page 11-13 (for the menu).

As the preceding table indicates, when you enable GVRP, a port that has a Tagged or Untagged static VLAN has the option for both generating advertisements and dynamically joining other VLANs.

Note

In table 11-4, above, the Unknown VLAN parameters are configured on a per-port basis using the CLI. The Tagged, Untagged, Auto, and Forbid options are configured per static VLAN on every port, using either the menu interface or the CLI.

Because dynamic VLANs operate as Tagged VLANs, and because a tagged port on one device cannot communicate with an untagged port on another device, HP recommends that you use Tagged VLANs for the static VLANs you will use to generate advertisements.

GVRP and VLAN Access Control

When you enable GVRP on a switch, the default GVRP parameter settings allow all of the switch's ports to transmit and receive dynamic VLAN advertisements (GVRP advertisements) and to dynamically join VLANs. The two preceding sections describe the per-port features you can use to control and limit VLAN propagation. To summarize, you can:

- Allow a port to advertise and/or join dynamic VLANs (Learn mode—the default).
- Allow a port to send VLAN advertisements, but not receive them from other devices; that is, the port cannot dynamically join a VLAN but other devices can dynamically join the VLANs it advertises (Block mode).
- Prevent a port from participating in GVRP operation (Disable mode).

Port-Leave From a Dynamic VLAN

A dynamic VLAN continues to exist on a port for as long as the port continues to receive advertisements of that VLAN from another device connected to that port or until you:

- Convert the VLAN to a static VLAN (See "Converting a Dynamic VLAN to a Static VLAN" on page 11-20.)
- Reconfigure the port to **Block** or **Disable**
- Disable GVRP
- Reboot the switch

The time-to-live for dynamic VLANs is 10 seconds. That is, if a port has not received an advertisement for an existing dynamic VLAN during the last 10 seconds, the port removes itself from that dynamic VLAN.

Planning for GVRP Operation

These steps outline the procedure for setting up dynamic VLANs for a segment.

1. Determine the VLAN topology you want for each segment (broadcast domain) on your network.
2. Determine the VLANs that must be static and the VLANs that can be dynamically propagated.
3. Determine the device or devices on which you must manually create static VLANs in order to propagate VLANs throughout the segment.
4. Determine security boundaries and how the individual ports in the segment will handle dynamic VLAN advertisements. (See table 11-3 on page 11-38 and table 11-4 on page 11-40.)
5. Enable GVRP on all devices you want to use with dynamic VLANs and configure the appropriate “Unknown VLAN” parameter (**Learn**, **Block**, or **Disable**) for each port.
6. Configure the static VLANs on the switch(es) where they are needed, along with the per-VLAN parameters (**Tagged**, **Untagged**, **Auto**, and **Forbid**—see table 11-4 on page 11-40) on each port.
7. Dynamic VLANs will then appear automatically, according to the configuration options you have chosen.
8. Convert dynamic VLANs to static VLANs where you want dynamic VLANs to become permanent.

Configuring GVRP On a Switch

The procedures in this section describe how to:

- View the GVRP configuration on a switch
- Enable and disable GVRP on a switch
- Specify how individual ports will handle advertisements

To view or configure static VLANs for GVRP operation, refer to “Per-Port Static VLAN Configuration Options” on page 11-7.

Menu: Viewing and Configuring GVRP

1. From the Main Menu, select:

2. Switch Configuration . . .
8. VLAN Menu . . .
1. VLAN Support

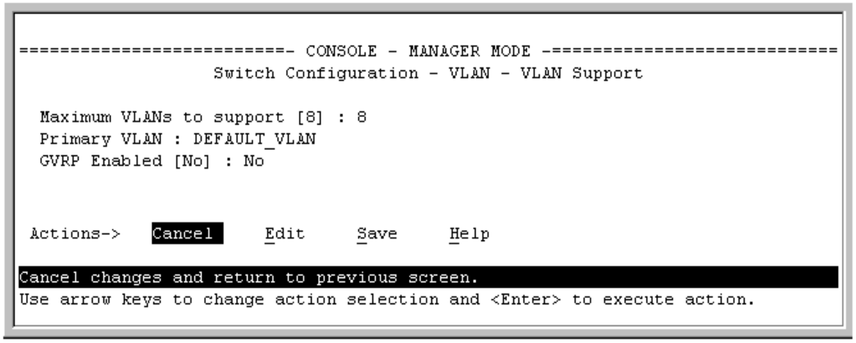

```
==========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - VLAN - VLAN Support

  Maximum VLANs to support [8] : 8
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled [No] : No


 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 11-27. The VLAN Support Screen (Default Configuration)

2. Do the following to enable GVRP and display the Unknown VLAN fields:
   a. Press E (for **Edit**).
   b. Use ↓ to move the cursor to the **GVRP Enabled** field.
   c. Press the Space bar to select **Yes**.
   d. Press ↓ again to display the **Unknown VLAN** fields.

The Unknown VLAN fields enable you to configure each port to:

- Learn - Dynamically join any advertised VLAN and advertise all VLANs learned through other ports.
- Block - Do not dynamically join any VLAN, but still advertise all VLANs learned through other ports.
- Disable - Ignore and drop all incoming advertisements and do not transmit any advertisements.

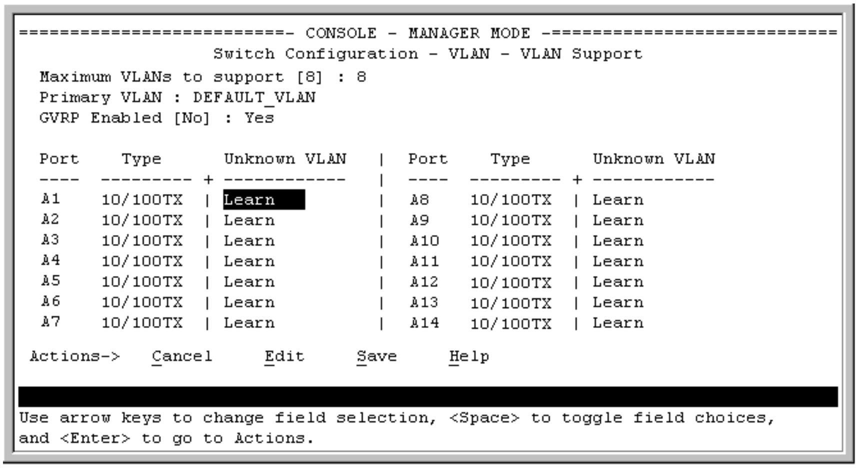

```
==========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - VLAN - VLAN Support
  Maximum VLANs to support [8] : 8
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled [No] : Yes

  Port   Type       Unknown VLAN   |  Port   Type       Unknown VLAN
  ----   --------- + ------------   |  ----   --------- + ------------
  A1     10/100TX  | Learn          |  A8     10/100TX  | Learn
  A2     10/100TX  | Learn          |  A9     10/100TX  | Learn
  A3     10/100TX  | Learn          |  A10    10/100TX  | Learn
  A4     10/100TX  | Learn          |  A11    10/100TX  | Learn
  A5     10/100TX  | Learn          |  A12    10/100TX  | Learn
  A6     10/100TX  | Learn          |  A13    10/100TX  | Learn
  A7     10/100TX  | Learn          |  A14    10/100TX  | Learn

 Actions->   Cancel     Edit     Save     Help


Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Figure 11-28. Example Showing Default Settings for Handling Advertisements

3. Use the arrow keys to select the port you want, and the Space bar to select Unknown VLAN option for any ports you want to change.

4. When you finish making configuration changes, press Enter, then S (for **Save**) to save your changes to the Startup-Config file.

CLI: Viewing and Configuring GVRP

GVRP Commands Used in This Section

show gvrp	below
gvrp	page 11-45
unknown-vlans	page 11-45

Displaying the Switch's Current GVRP Configuration. This command shows whether GVRP is disabled, along with the current settings for the maximum number of VLANs and the current Primary VLAN. (For more on the last two parameters, see "Port-Based Virtual LANs (Static VLANs)" on page 11-3.)

Syntax: show gvrp Shows the current settings.

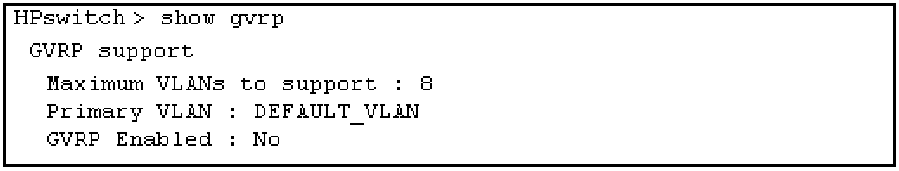

```
HPswitch > show gvrp
 GVRP support
  Maximum VLANs to support : 8
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled : No
```

Figure 11-29. Example of "Show GVRP" Listing with GVRP Disabled

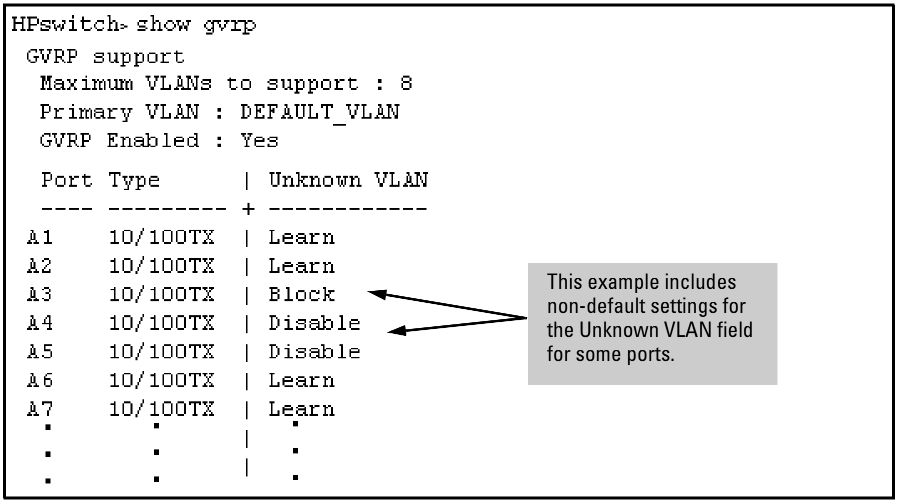

```
HPswitch> show gvrp
 GVRP support
  Maximum VLANs to support : 8
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled : Yes

  Port Type      | Unknown VLAN
  ---- --------- + ------------
 A1    10/100TX  | Learn
 A2    10/100TX  | Learn
 A3    10/100TX  | Block
 A4    10/100TX  | Disable
 A5    10/100TX  | Disable
 A6    10/100TX  | Learn
 A7    10/100TX  | Learn
  .       .      |   .
  .       .      |   .
  .       .      |   .
```

This example includes non-default settings for the Unknown VLAN field for some ports.

Figure 11-30. Example of Show GVRP Listing with GVRP Enabled

Enabling and Disabling GVRP on the Switch. This command enables GVRP on the switch.

Syntax: gvrp

This example enables GVRP:

```
HPswitch(config)# gvrp
```

This example disables GVRP operation on the switch:

```
HPswitch(config)# no gvrp
```

Enabling and Disabling GVRP On Individual Ports. When GVRP is enabled on the switch, use the **unknown-vlans** command to change the Unknown VLAN field for one or more ports. You can use this command at either the Manager level or the interface context level for the desired port(s).

Syntax: interface <*port-list*> unknown-vlans <learn | block | disable>

Changes the Unknown VLAN field setting for the specified port(s).

For example, to change and view the configuration for ports A1-A2 to **Block**:

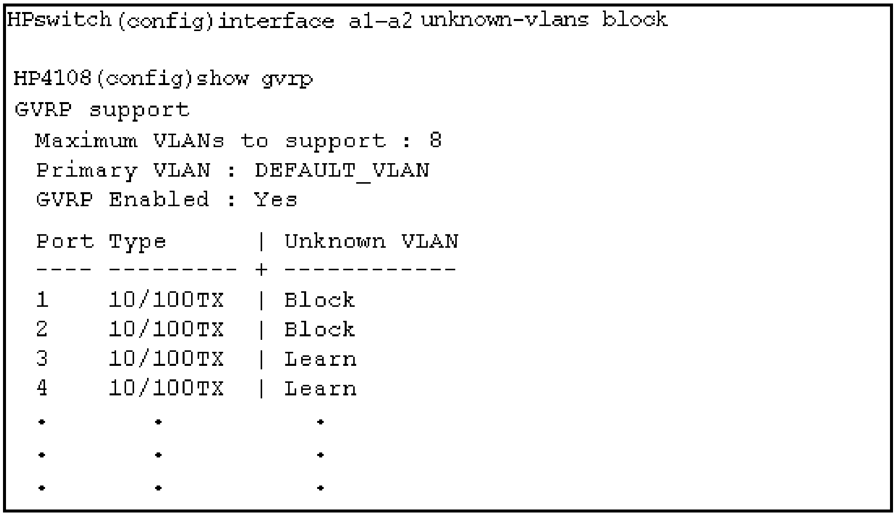

```
HPswitch(config)interface a1-a2 unknown-vlans block

HP4108(config)show gvrp
GVRP support
  Maximum VLANs to support : 8
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled : Yes

  Port Type      | Unknown VLAN
  ---- --------- + ------------
  1    10/100TX  | Block
  2    10/100TX  | Block
  3    10/100TX  | Learn
  4    10/100TX  | Learn
  .       .          .
  .       .          .
  .       .          .
```

Displaying the Static and Dynamic VLANs Active on the Switch. The **show vlans** command lists all VLANs present in the switch.

Syntax: show vlans

For example, in the following illustration, switch "B" has one static VLAN (the default VLAN), with GVRP enabled and port 1 configured to **Learn** for Unknown VLANs. Switch "A" has GVRP enabled and has three static VLANs: the default VLAN, VLAN-222, and VLAN-333. In this scenario, switch B will dynamically join VLAN-222 and VLAN-333:

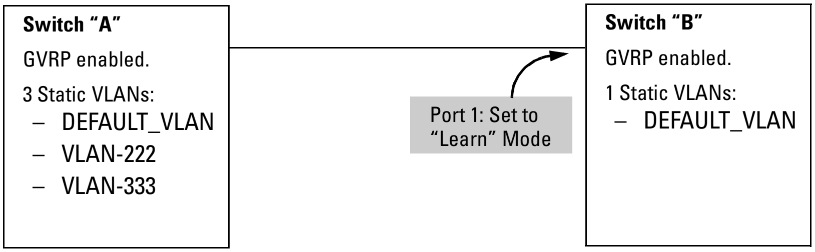

The **show vlans** command lists the dynamic (and static) VLANs in switch "B" after it has learned and joined VLAN-222 and VLAN-333.

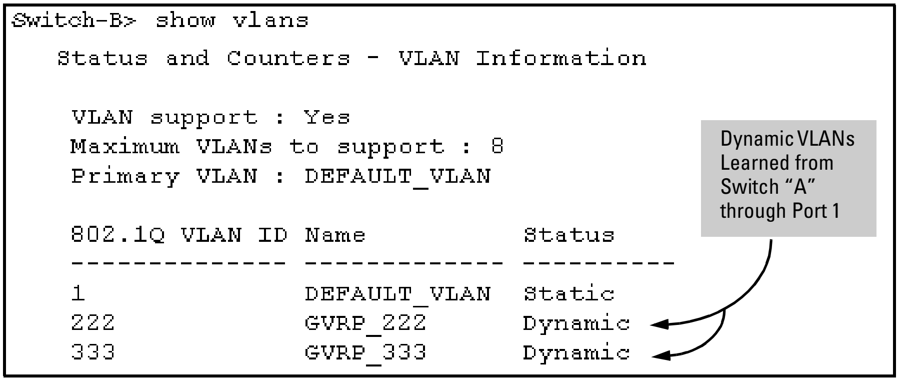

```
Switch-B> show vlans
   Status and Counters - VLAN Information

    VLAN support : Yes
    Maximum VLANs to support : 8
    Primary VLAN : DEFAULT_VLAN

    802.1Q VLAN ID Name          Status
    -------------- ------------- ----------
    1              DEFAULT_VLAN  Static
    222            GVRP_222      Dynamic
    333            GVRP_333      Dynamic
```

Dynamic VLANs Learned from Switch "A" through Port 1

Figure 11-31. Example of Listing Showing Dynamic VLANs

Converting a Dynamic VLAN to a Static VLAN. If a port on the switch has joined a dynamic VLAN, you can use the following command to convert that dynamic VLAN to a static VLAN:

Syntax: static <*dynamic-vlan-id*>

For example, to convert dynamic VLAN 333 (from the previous example) to a static VLAN:

```
HPswitch(config)# static 333
```

When you convert a dynamic VLAN to a static VLAN, all ports on the switch are assigned to the VLAN in Auto mode.

Web: Viewing and Configuring GVRP

To view, enable, disable, or reconfigure GVRP:

1. Click on the **Configuration** tab.
2. Click on [VLAN Configuration] and do the following:
   - To enable or disable GVRP, click on **GVRP Enabled**.
   - To change the Unknown VLAN field for any port:
     i. Click on [GVRP Security] and make the desired changes.
     ii. Click on [Apply] to save and implement your changes to the Unknown VLAN fields.

For web-based Help on how to use the web browser interface screen, click on the [?] button provided on the web browser screen.

GVRP Operating Notes

- A dynamic VLAN must be converted to a static VLAN before it can have an IP address.
- The total number of VLANs on the switch (static and dynamic combined) cannot exceed the current Maximum VLANs setting. For example, in the factory default state, the switch supports eight VLANs. Thus, in a case where four static VLANs are configured on the switch, the switch can accept up to four additional VLANs in any combination of static and dynamic. Any additional VLANs advertised to the switch will not be added unless you first increase the Maximum VLANs setting. In the Menu interface, click on **2. Switch Configuration ... | 8. VLAN Menu | 1. VLAN Support**. In the global config level of the CLI, use **max-vlans**.

- Converting a dynamic VLAN to a static VLAN and then executing the **write memory** command saves the VLAN in the startup-config file and makes it a permanent part of the switch's VLAN configuration.
- Within the same broadcast domain, a dynamic VLAN can pass through a device that is not GVRP-aware. This is because a hub or a switch that is not GVRP-ware will flood the GVRP (multicast) advertisement packets out all ports.
- GVRP assigns dynamic VLANs as Tagged VLANs. To configure the VLAN as Untagged, you must first convert it to a static VLAN.
- Rebooting a switch on which a dynamic VLAN exists deletes that VLAN. However, the dynamic VLAN re-appears after the reboot if GVRP is enabled and the switch again receives advertisements for that VLAN through a port configured to add dynamic VLANs.
- By receiving advertisements from other devices running GVRP, the switch learns of static VLANs on those other devices and dynamically (automatically) creates tagged VLANs on the links to the advertising devices. Similarly, the switch advertises its static VLANs to other GVRP-aware devices, as well as the dynamic VLANs the switch has learned.
- A GVRP-enabled switch does not advertise any GVRP-learned VLANs out of the port(s) on which it originally learned of those VLANs.

Multimedia Traffic Control with IP Multicast (IGMP)

Contents

Overview

IGMP Features

Feature	Default	Menu	CLI	Web
view igmp configuration	n/a	—	page 12-6	—
show igmp status for multicast groups used by the selected VLAN	n/a	—	Yes	—
enabling or disabling IGMP (Requires VLAN ID Context)	disabled	—	page 12-7	page 12-9
per-port packet control	auto	—	page 12-8	—
IGMP traffic priority	normal	—	page 12-8	—
querier	enabled	—	page 12-9	—

Without IGMP enabled, the switch simply floods all IP multicast traffic it receives on a given VLAN through all ports on that VLAN (except the port on which it received the traffic). This can result in significant and unnecessary bandwidth usage in networks where IP multicast traffic is a factor. Data-Driven IGMP reduces this problem by authorizing the switch to restrict multicast traffic only to ports where a given multicast group should flow. Series 5300XL switches use data-driven IGMP to better control IP multicast traffic.

In a network where IP multicast traffic is transmitted for various multimedia applications, you can use IGMP on the switch to reduce unnecessary bandwidth usage on a per-port basis. In the factory default state (IGMP disabled), the switch forwards all IGMP traffic to all ports, which can cause unnecessary bandwidth usage on ports not belonging to multicast groups. Enabling IGMP allows the ports to detect IGMP queries and report packets and manage IP multicast traffic through the switch.

IGMP is useful in multimedia applications such as LAN TV, desktop conferencing, and collaborative computing, where there is multipoint communication; that is, communication from one to many hosts, or communication originating from many hosts and destined for many other hosts. In such multipoint applications, IGMP will be configured on the hosts, and multicast traffic will be generated by one or more servers (inside or outside of the local network). Switches in the network (that support IGMP) can then be configured to direct the multicast traffic to only the ports where needed. If multiple VLANs are configured, you can configure IGMP on a per-VLAN basis.

Enabling IGMP allows the ports to detect IGMP queries and report packets and manage IP multicast traffic through the switch. If no other querier is detected, the switch will then also function as the querier. (If you need to disable the querier feature, you can do so by using the **no ip igmp querier** command. Refer to "Configuring the Querier Capability" on page 12-9.)

Note

IGMP configuration on the Series 5300XL switches operates at the VLAN context level. If you are not using VLANs, then configure IGMP in VLAN 1 (the default VLAN) context.

Terminology

- **IGMP:** Internet Group Management Protocol. This is an internal protocol of the Internet Protocol (IP) suite. IP manages multicast traffic by using switches, multicast routers, and hosts that support IGMP.
- **IGMP Device:** A switch or router running IGMP traffic control features.
- **IGMP Host:** An end-node device running an IGMP (multipoint, or multicast communication) application.
- **Leave Group:** A message sent by a host to the querier to indicate that the host has ceased to be a member of a specific multicast group.
- **Multicast Group:** A set of hosts, routers, and/or switches that send or receive multicast data streams to or from the same source(s). All devices in the group use the same multicast group address.
- **Report:** A message sent by a host to the querier to indicate that the host wants to be or is a member of a given group indicated in the report message.
- **Query:** A message sent from the querier (multicast router or switch) asking for a response from each host belonging to the multicast group. For more information, refer to "More on IGMP Operation" on page 12-10.
- **Querier:** A required IGMP device that facilitates the IGMP protocol and traffic flow on a given LAN. This device tracks which ports are connected to devices (IGMP clients) that belong to specific multicast

groups, and triggers updates of this information. With IGMP enabled, the Series 5300XL switches use data from the Querier to determine whether to forward or block multicast traffic on specific ports. When the switch has an IP address on a given VLAN, it automatically operates as a Querier for that VLAN if it does not detect a multicast router or another switch functioning as a Querier. For more information, refer to "Querier Operation" on page 12-20.

IGMP Operating Features

In the factory default configuration, IGMP is disabled. If multiple VLANs are not configured, you must configure IGMP on the default VLAN (DEFAULT_VLAN; VID = 1). If multiple VLANs are configured, you must configure IGMP on a per-VLAN basis for each VLAN in which you want IGMP to operate'. When you use either the CLI or the web browser interface to enable IGMP on the switch or a VLAN, the switch forwards IGMP traffic only to ports belonging to multicast groups. Using the console enables these additional options:

- **Forward with High Priority.** Disabling this parameter (the default) causes the switch or VLAN to process IP multicast traffic, along with other traffic, in the order received (normal priority). Enabling this parameter causes the switch or VLAN to give a higher priority to IP multicast traffic than to other traffic.
- **Auto/Blocked/Forward:** You can use the console to configure individual ports to any of the following states:
  - **Auto** (the default): Causes the switch to interpret IGMP packets and to filter IP multicast traffic based on the IGMP packet information for ports belonging to a multicast group. This means that IGMP traffic will be forwarded on a specific port only if an IGMP host or multicast router is connected to the port.
  - **Forward:** Causes the switch to forward all IGMP and IP multicast transmissions through the port.
  - **Blocked:** Causes the switch to drop all IGMP transmissions received from a specific port and to block all outgoing IP Multicast packets for that port. This has the effect of preventing IGMP traffic from moving through specific ports.

Note

If you configure an individual port as blocked, but that port has either a **forward** configured on another VLAN or belongs to another VLAN that is connected to a router, blocking will not occur.

- **Querier:** In the default state (enabled), eliminates the need for a multicast router. In most cases, HP recommends that you leave this parameter in the default "enabled" state even if you have a multicast router performing the querier function in your multicast group. Refer to "Querier Operation" on page 12-20.
- **IP Addressing:** You can configure IGMP on a VLAN regardless of whether the VLAN has an IP address. The limitation imposed by a lack of an IP address on a VLAN is that the switch cannot operate as a querier on that VLAN. Refer to "IGMP Operates With or Without IP Addressing" on page 12-15.
- **Fast-Leave IGMP:** This automatic feature improves performance by accelerating blocking on a port where a client has left a multicast group and the switch does not detect any other end nodes on the port. Refer to "Fast-Leave IGMP" on page 12-16.
- **Forced Fast-Leave IGMP:** Enables you to configure IGMP to speed up the process of blocking unnecessary IGMP traffic to a switch port connected to multiple end nodes. Refer to "Forced Fast-Leave IGMP" on page 12-18.

Note

Whenever IGMP is enabled, the switch generates an Event Log message indicating whether querier functionality is enabled.

For more information, refer to "More on IGMP Operation" on page 12-10.

CLI: Configuring and Displaying IGMP

IGMP Commands	Page
show ip igmp configuration	12-6
ip igmp	12-7
high-priority-forward	12-8
auto <[ethernet] *<port-list>*	12-8
blocked <[ethernet] *<port-list>*	12-8
forward <[ethernet] *<port-list>*	12-8
querier	12-9
show ip igmp	See "Internet Group Mulitcast Protocol (IGMP) Status" on page B-18

Viewing the Current IGMP Configuration. This command lists the IGMP configuration for all VLANs configured on the switch or for a specific VLAN.

Syntax: show ip igmp config — IGMP configuration for all VLANs on the switch

show ip igmp < *vlan-id* > config — IGMP configuration for a specific VLAN on the switch, including per-port data

(For IGMP operating status, see "Internet Group Management Protocol (IGMP) Status" on page B-18.)

For example, suppose you have the following VLAN and IGMP configurations on the switch:

VLAN ID	VLAN Name	IGMP Enabled	Forward with High Priority	Querier
1	DEFAULT_VLAN	Yes	No	No
22	VLAN-2	Yes	Yes	Yes
33	VLAN-3	No	No	No

You could use the CLI to display this data as follows:

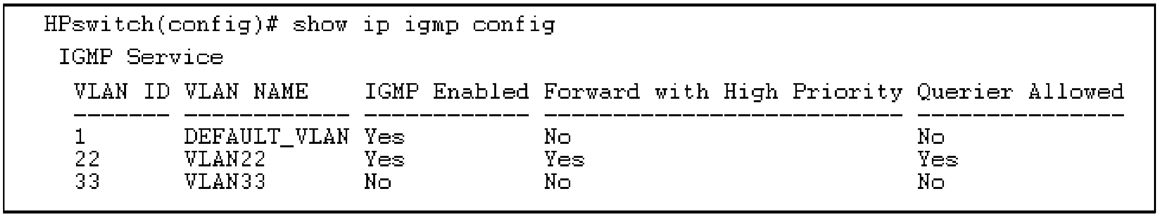

```
HPswitch(config)# show ip igmp config
 IGMP Service
  VLAN ID VLAN NAME    IGMP Enabled Forward with High Priority Querier Allowed
  ------- ------------ ------------ -------------------------- ---------------
  1       DEFAULT_VLAN Yes          No                         No
  22      VLAN22       Yes          Yes                        Yes
  33      VLAN33       No           No                         No
```

Figure 12-1. Example Listing of IGMP Configuration for All VLANs in the Switch

The following version of the show ip igmp command includes the VLAN ID (*VID*) designation, and combines the above data with the IGMP per-port configuration:

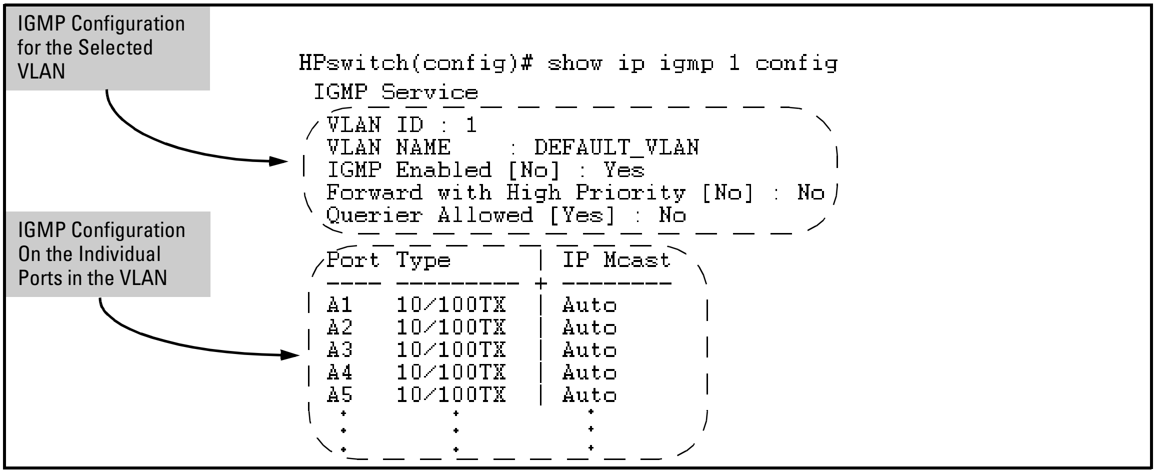

IGMP Configuration for the Selected VLAN

IGMP Configuration On the Individual Ports in the VLAN

```
HPswitch(config)# show ip igmp 1 config
 IGMP Service
 VLAN ID : 1
 VLAN NAME     : DEFAULT_VLAN
 IGMP Enabled [No] : Yes
 Forward with High Priority [No] : No
 Querier Allowed [Yes] : No

 Port Type       | IP Mcast
 ---- --------- + --------
 A1   10/100TX  | Auto
 A2   10/100TX  | Auto
 A3   10/100TX  | Auto
 A4   10/100TX  | Auto
 A5   10/100TX  | Auto
  .       .          .
  .       .          .
  .       .          .
```

Figure 12-2. Example Listing of IGMP Configuration for A Specific VLAN

Enabling or Disabling IGMP on a VLAN. You can enable IGMP on a VLAN, along with the last-saved or default IGMP configuration (whichever was most recently set), or you can disable IGMP on a selected VLAN. Note that this command must be executed in a VLAN context.

Syntax: [no] ip igmp

For example, here are methods to enable and disable IGMP on the default VLAN (VID = 1).

`HPswitch(config)# vlan 1 ip igmp`	Enables IGMP on VLAN 1.
`HPswitch(vlan-1)# ip igmp`	Same as above.
`HPswitch(config)# no vlan 1 ip igmp`	Disables IGMP on VLAN 1.

Note

If you disable IGMP on a VLAN and then later re-enable IGMP on that VLAN, the switch restores the last-saved IGMP configuration for that VLAN. For more on how switch memory operates, see Chapter 5, “Switch Memory and Configuration”.

You can also combine the **ip igmp** command with other IGMP-related commands, as described in the following sections.

Configuring Per-Port IGMP Packet Control. Use this command in the VLAN context to specify how each port should handle IGMP traffic.

Syntax: vlan < *vlan-id* > ip igmp
[auto <*port-list*> | blocked <*port-list*> | forward <*port-list*>]

Default: auto

Note

The blocked feature does not operate on a port if the port either detects a router or belongs to multiple VLANs and is configured with **forward** on one of these other VLANs.

For example, suppose you wanted to configure IGMP as follows for VLAN 1 on the 10/100 ports on the switch:

Ports A1-A7	auto	Filter multicast traffic. Forward IGMP traffic to hosts on these ports that belong to the multicast group for which the traffic is intended. (Also forward any multicast traffic through any of these ports that is connected to a multicast router.)
Port A8	forward	Forward all multicast traffic through this port.
Ports A9-A12	blocked	Drop all multicast traffic received from devices on these ports, and prevent any outgoing multicast traffic from moving through these ports.

Depending on the privilege level, you could use one of the following commands to configure IGMP on VLAN 1 with the above settings:

```
HPswitch(config)# vlan 1 ip igmp auto a1-a7 forward a8 blocked
a9-a12
```

```
HPswitch(vlan-1)# ip igmp auto a1-a7 forward a8 blocked a9-
a12
```

The following command displays the VLAN and per-port configuration resulting from the above commands.

```
HPswitch> show ip igmp 1 config
```

Configuring IGMP Traffic Priority. This command assigns “high” priority to IGMP traffic or returns a high-priority setting to “normal” priority.

Syntax: vlan < *vlan-id* > ip igmp high-priority-forward

Default: normal

```
HPswitch(config)# vlan 1 ip igmp
            high-priority-forward
```

Configures high priority for IGMP traffic on VLAN 1.

```
HPswitch(vlan-1)# ip igmp
           high-priority-forward
```

Same as above command, but in the VLAN 1 context level.

```
HPswitch(vlan 1)# no ip igmp
           high-priority-forward
```

Returns IGMP traffic to "normal" priority.

```
HPswitch> show ip igmp config
```

Show command to display results of above high-priority commands.

Configuring the Querier Capability. The default querier capability is "enabled". This means that the switch will be allowed to operate as the querier on a VLAN (but does not force it to do so). This command disables or re-enables the querier capability. (Refer also to "General IGMP Operation" on page 12-11.)

Syntax: [no] vlan < *vlan-id* > ip igmp querier

Default: Yes

```
HPswitch(config)# no vlan 1 ip
                 igmp querier
```

Disables the querier capability on VLAN 1.

```
HPswitch> show ip igmp config
```

Show command to display results of above querier command.

Web: Enabling or Disabling IGMP

In the web browser interface you can enable or disable IGMP on a per-VLAN basis. To configure other IGMP features, telnet to the switch console and use the CLI.

To Enable or Disable IGMP

1. Click on the **Configuration** tab.
2. Click on Device Features.
3. If more than one VLAN is configured, use the VLAN pull-down menu to select the VLAN on which you want to enable or disable IGMP.
4. Use the Multicast Filtering (IGMP) menu to enable or disable IGMP.
5. Click on Apply Changes to implement the configuration change.

For web-based help on how to use the web browser interface screen, click on the [?] button provided on the web browser screen.

More on IGMP Operation

Multicast traffic management uses switches, multicast routers, and hosts that support IGMP. (In Hewlett-Packard's implementation of IGMP, a multicast router is not necessary as long as a switch is configured to support IGMP with the **querier** feature enabled.) The multicast group running version 2 of IGMP uses three fundamental types of messages to communicate:

- **Query:** A message sent from the querier (multicast router or switch) asking for a response from each host belonging to the multicast group. If a multicast router supporting IGMP is not present, then the switch must assume this function in order to elicit group membership information from the hosts on the network. (If you need to disable the querier feature, you can do so through the CLI, using the IGMP configuration MIB. Refer to "Configuring the Querier Capability" on page 12-9.)
- **Report (Join):** A message sent by a host to the querier to indicate that the host wants to be or is a member of a given group indicated in the report message.
- **Leave Group:** A message sent by a host to the querier to indicate that the host has ceased to be a member of a specific multicast group.

Thus, IGMP identifies members of a multicast group (within a subnet) and allows IGMP-configured hosts (and routers) to join or leave multicast groups.

IGMP Data. To display data showing active group addresses, reports, queries, querier access port, and active group address data (port, type, and access), see "IP Multicast (IGMP) Status" on page B-18.

General IGMP Operation

With IGMP enabled, the switch examines the IGMP packets it receives to learn which of its ports are linked to IGMP hosts and multicast routers/queriers belonging to any multicast group. The switch becomes a querier if it does not discover a multicast router/querier on the network.

Once the switch learns the port location of the hosts belonging to any particular multicast group, it can direct group traffic to only those ports, resulting in bandwidth savings on ports where group members do not reside.

The following example illustrates this operation.

Figure 12-3 on page 12-12 shows a network running IGMP.

- PCs 1 and 4, switch 2, and all of the routers are members of an IP multicast group. (The routers operate as queriers.)
- Switch 1 ignores IGMP traffic and does not distinguish between IP multicast group members and non-members. Thus, it is sending large amounts of unwanted multicast traffic out the ports to PCs 2 and 3.
- Switch 2 is recognizing IGMP traffic and learns that PC 4 is in the IP multicast group receiving multicast data from the video server (PC X). Switch 2 then sends the multicast data only to the port for PC 4, thus avoiding unwanted multicast traffic on the ports for PCs 5 and 6.

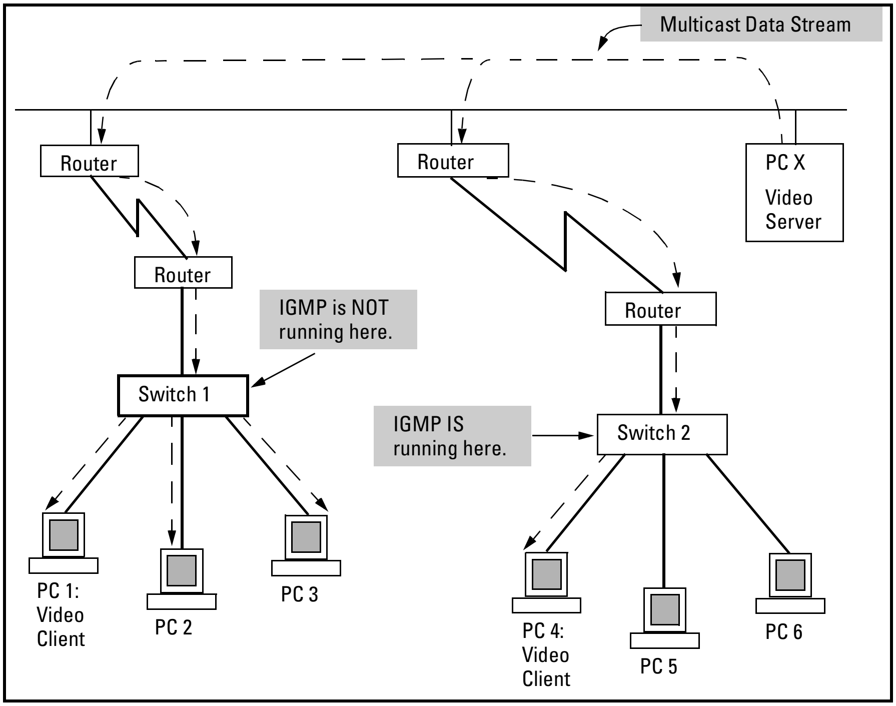

Figure 12-3. The Advantage of Using IGMP

The next figure (12-4) shows a network running IP multicasting using IGMP without a multicast router. In this case, the IGMP-configured switch runs as a querier.

PCs 2, 5, and 6 are members of the same IP multicast group.

IGMP is configured on switches 3 and 4. Either of these switches can operate as querier because a multicast router is not present on the network. (If an IGMP switch does not detect a querier, it automatically assumes this role, assuming the querier feature is enabled—the default—within IGMP, and that the switch has an IP address on the IGMP-enabled VLAN.)

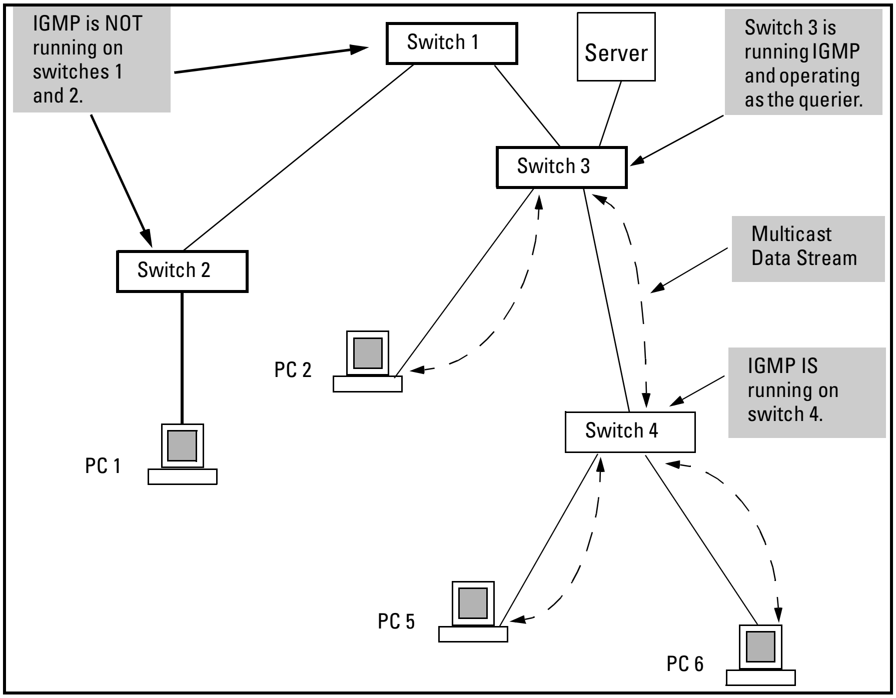
IGMP is NOT running on switches 1 and 2.
Switch 1
Server
Switch 3 is running IGMP and operating as the querier.
Switch 3
Multicast Data Stream
Switch 2
PC 2
IGMP IS running on switch 4.
Switch 4
PC 1
PC 5
PC 6

Figure 12-4. Isolating IP Multicast Traffic in a Network

- In the above figure, the multicast group traffic does not go to switch 1 and beyond because either the port on switch 3 that connects to switch 1 has been configured as blocked (through IGMP or another means) or there are no hosts connected to switch 1 or switch 2 that belong to the multicast group.
- For PC 1 to become a member of the same multicast group without flooding IP multicast traffic on all ports of switches 1 and 2, IGMP must be configured on both switches 1 and 2, and the port on Switch 3 that connects to Switch 1 must be unblocked. (Blocking can be due to the current IGMP configuration or to another reason unrelated to IGMP.)

Data-Driven IGMP

An IP multicast packet includes the multicast group (address) to which the packet belongs. When an IGMP client connected to a switch port needs to receive multicast traffic from a specific group, it joins the group by sending an IGMP report (join request) through the switch to the network. (The multicast group specified in the join request is determined by the requesting application running on the IGMP client.) When a networking device with IGMP enabled receives the join request for a specific group, it forwards any IP multicast traffic it receives for that group through the port on which the join request was received. To reduce unnecessary traffic, the networking device does not forward a given group's multicast packets to ports from which a join request for that group has not been received. (If the switch or router has not received any join requests for a given multicast group, it drops the traffic it receives for that group.)

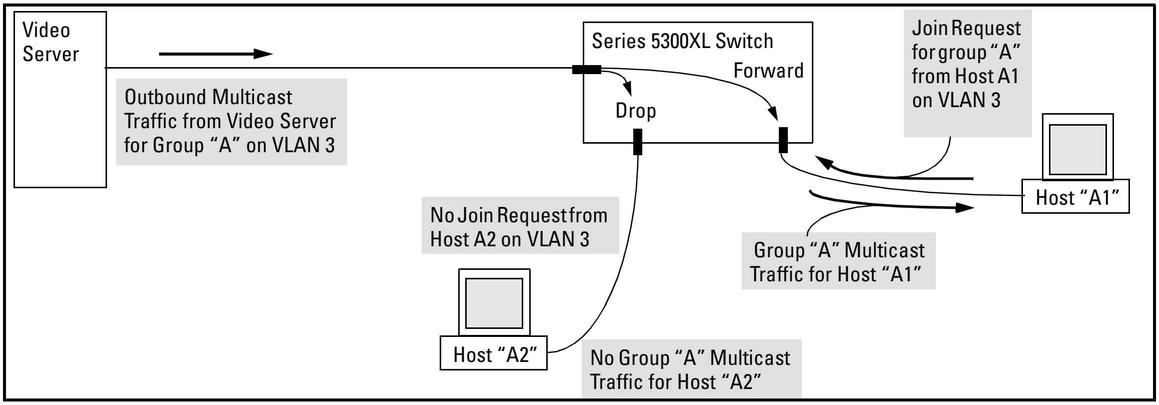
Video Server
Outbound Multicast Traffic from Video Server for Group "A" on VLAN 3
Series 5300XL Switch
Forward
Drop
Join Request for group "A" from Host A1 on VLAN 3
Host "A1"
No Join Request from Host A2 on VLAN 3
Group "A" Multicast Traffic for Host "A1"
Host "A2"
No Group "A" Multicast Traffic for Host "A2"

Figure 12-5. Example of Data-Driven IGMP Operation

Thus, after you enable IGMP on a VLAN configured in the switch, it continually listens for IGMP messages and IP multicast traffic on all ports in the VLAN, and forwards IGMP traffic for a given multicast address only through the port(s) on that VLAN where an IGMP report (join request) for that address was received from an IGMP client device.

Note

IP multicast traffic groups are identified by IP addresses in the range of 224.0.0.0 to 239.255.255.255.

Incoming IGMP packets intended for reserved, or "well-known" multicast addresses automatically flood through all ports (except the port on which the packets entered the switch). For more on this topic, see "Excluding Well-Known or Reserved Multicast Addresses from IP Multicast Filtering" on page 12-21.

Number of IP Multicast Addresses Allowed

The total of static multicast filters and IGMP multicast filters (addresses) together can range from 389 to 420, depending on the current **max-vlans** setting in the switch. If multiple VLANs are configured, then each filter is counted once per VLAN in which it is used.

Table 12-1. Maximum Allowed Number of Multicast Filters

Max-VLANs Setting	Maximum # of Multicast Filters (Static and IGMP Combined)
1 (the minimum)	420
8 (the default)	413
32 or higher	389

IGMP Operates With or Without IP Addressing

You can configure IGMP on VLANs that do not have IP addressing. The benefit of IGMP without IP addressing is a reduction in the number of IP addresses you have to use and configure. This can be significant in a network with a large number of VLANs. The limitation on IGMP without IP addressing is that the switch cannot become Querier on any VLANs for which it has no IP address—so the network administrator must ensure that another IGMP device will act as Querier and that an additional IGMP device is available as a backup Querier.

IGMP Function Available With IP Addressing Configured on the VLAN	Available *Without* IP Addressing?	Operating Differences Without an IP Address
Drop multicast group traffic for which there have been no join requests from IGMP clients connected to ports on the VLAN.	Yes	None
Forward multicast group traffic to any port on the VLAN that has received a join request for that multicast group.	Yes	None
Forward join requests (reports) to the Querier.	Yes	None
Configure individual ports in the VLAN to **Auto** (the default)/**Blocked**, or **Forward**.	Yes	None
Configure IGMP traffic forwarding to normal or high-priority forwarding.	Yes	None

IGMP Function Available With IP Addressing Configured on the VLAN	Available *Without* IP Addressing?	Operating Differences Without an IP Address
Age-Out IGMP group addresses when the last IGMP client on a port in the VLAN leaves the group.	Yes	Requires that another IGMP device in the VLAN has an IP address and can operate as Querier. This can be a multicast router or another switch configured for IGMP operation. (HP recommends that the VLAN also include a device operating as a backup Querier in case the device operating as the primary Querier fails for any reason.
Support Fast-Leave IGMP (below) and Forced Fast-Leave IGMP (page 18).	Yes	
Support automatic Querier election.	No	Querier operation not available.
Operate as the Querier.	No	Querier operation not available.
Provide a backup Querier.	No	Querier operation not available.

Fast-Leave IGMP

IGMP Operation Presents a "Delayed Leave" Problem. Where multiple IGMP clients are connected to the same port on an IGMP device (switch or router), if only one IGMP client joins a given multicast group, then later sends a Leave Group message and ceases to belong to that group, the IGMP device retains that IGMP client in its IGMP table and continues forwarding IGMP traffic to the IGMP client until the Querier triggers confirmation that no other group members exist on the same port. This means that the switch continues to transmit unnecessary multicast traffic through the port until the Querier renews multicast group status.

Fast-Leave IGMP Reduces Leave Delays. Fast-Leave IGMP operates on a port if an IGMP client connects to the port and there are no other end nodes detected on that port. In this case, when the client leaves a multicast group, Fast-Leave IGMP automatically accelerates the blocking of further, unnecessary multicast traffic from that group to the former IGMP client. This improves performance by reducing the amount of multicast traffic going through the port to the IGMP client after the client leaves a multicast group. IGMP in the Series 5300XL switches automatically uses this Fast-Leave feature.

Automatic Fast-Leave Operation. If a Series 5300XL switch port is:

a. Connected to only one end node

b. The end node currently belongs to a multicast group; i.e. is an IGMP client

c. The end node subsequently leaves the multicast group

Then the switch does not need to wait for the Querier status update interval, but instead immediately removes the IGMP client from its IGMP table and ceases transmitting IGMP traffic to the client. (If the switch detects multiple end nodes on the port, automatic Fast-Leave does not activate—regardless of whether one or more of these end nodes are IGMP clients.)

In the next figure, automatic Fast-Leave operates on the switch ports for IGMP clients "3A" and "5B", but not on the switch port for IGMP clients "7A" and 7B, Server "7C", and printer "7D".

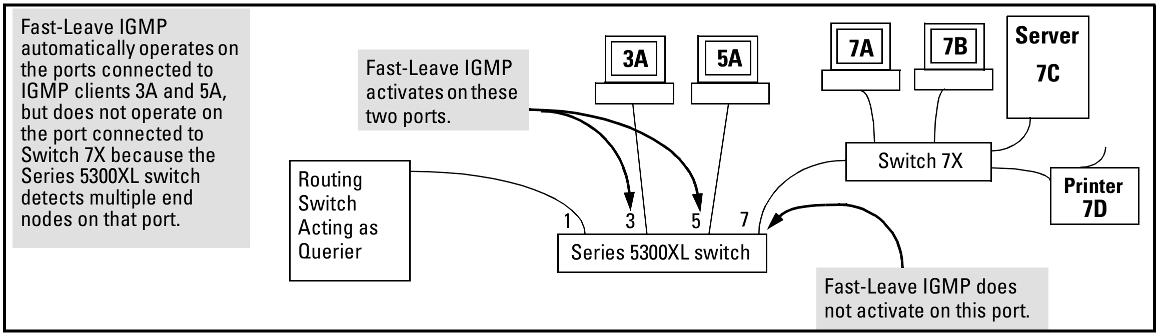
Fast-Leave IGMP automatically operates on the ports connected to IGMP clients 3A and 5A, but does not operate on the port connected to Switch 7X because the Series 5300XL switch detects multiple end nodes on that port.
Fast-Leave IGMP activates on these two ports.
3A
5A
7A
7B
Server 7C
Routing Switch Acting as Querier
Switch 7X
Printer 7D
1
3
5
7
Series 5300XL switch
Fast-Leave IGMP does not activate on this port.

Figure 12-6. Example of Automatic Fast-Leave IGMP Criteria

When client "3A" running IGMP is ready to leave the multicast group, it transmits a Leave Group message. Because the switch knows that there is only one end node on port 3, it removes the client from its IGMP table and halts multicast traffic (for that group) to port 3. If the switch is not the Querier, it does not wait for the actual Querier to verify that there are no other group members on port 3. If the switch itself is the Querier, it does not query port 3 for the presence of other group members.

Note that Fast-Leave operation does not distinguish between end nodes on the same port that belong to different VLANs. Thus, for example, even if all of the devices on port 7 in figure 12-6 belong to different VLANs, Fast-Leave does not operate on port 7.

Forced Fast-Leave IGMP

Forced Fast-Leave IGMP Features

Feature	Default	Menu	CLI	Web
view the Forced Fast-Leave configuration				
view the switch's Forced Fast-Leave state	n/a	—	page 12-19	—
configure Forced Fast-Leave				
configure Forced Fast-Leave for an individual port	2 (disabled)	—	page 12-20	—

Forced Fast-Leave IGMP speeds up the process of blocking unnecessary IGMP traffic to a switch port that is connected to multiple end nodes. (This feature does not activate on ports where the switch detects only one end node. Instead, the regular Fast Leave described in the preceding section activates.) For example, in figure 12-6, even if you configured Forced Fast-Leave on all ports in the switch, the feature would activate only on port 7 (which has multiple end nodes) when a Leave Group request arrived on that port.

When a port having multiple end nodes receives a Leave Group request from one end node for a given multicast group "X", Forced Fast-Leave activates and waits a small amount of time to receive a join request from any other group "X" member on that port. If the port does not receive a join request for that group within the forced-leave interval, the switch then blocks any further group "X" traffic to the port.

Configuration Options for Forced Fast-Leave

Feature	Default	Settings	Function
Forced Fast-Leave state	**2** (disabled)	**1** (enabled) **2** (disabled)	Uses the **setmib** command to enable or disable Forced Fast-Leave on individual ports. When enabled on a port, Forced Fast-Leave operates only if the switch detects multiple end nodes (and at least one IGMP client) on that port.

Listing the Forced Fast-Leave Configuration. The Forced Fast-Leave configuration includes the state (enabled or disabled) for each port and the Forced-Leave Interval for all ports on the switch.

To list the Forced Fast-Leave state for all ports in the switch:

Syntax:

```
HPswitch# walkmib hpSwitchIgmpPortForcedLeaveState.1
or
HPswitch# walkmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5.1
```

For example:

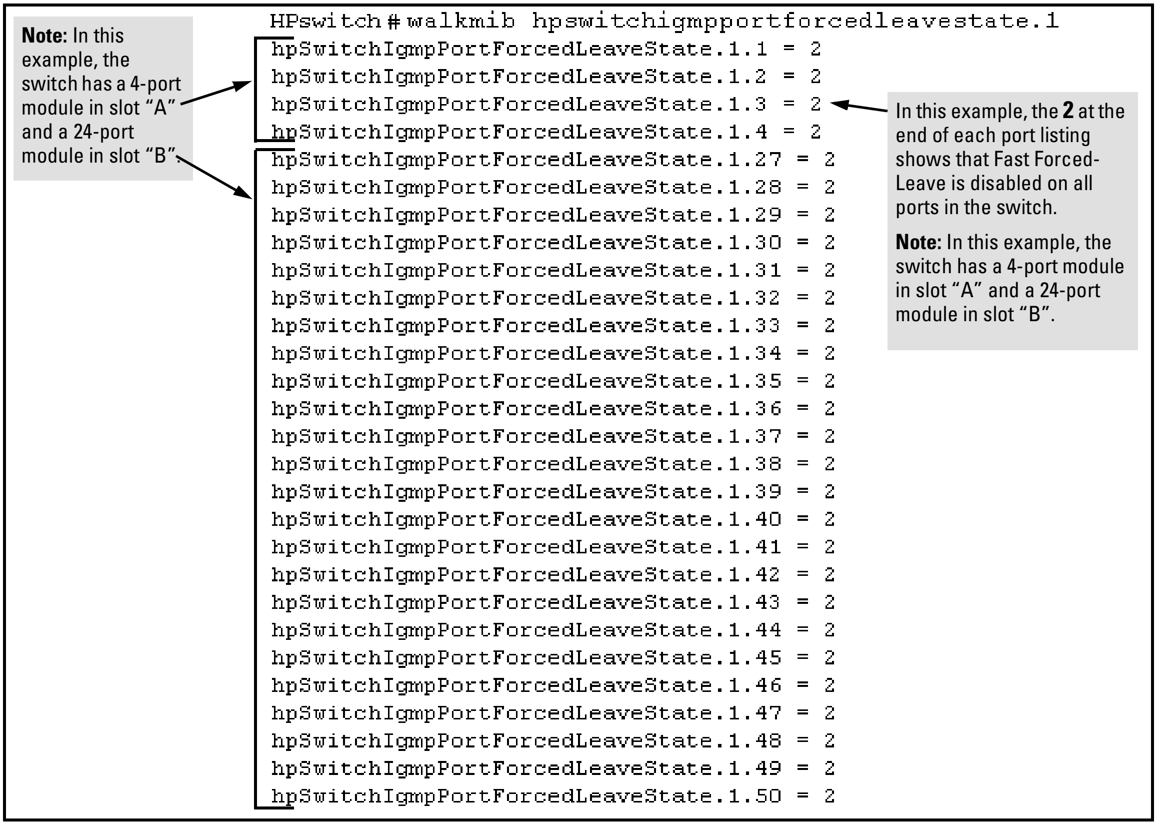

Note: In this example, the switch has a 4-port module in slot “A” and a 24-port module in slot “B”.

```
HPswitch # walkmib hpswitchigmpportforcedleavestate.1
hpSwitchIgmpPortForcedLeaveState.1.1 = 2
hpSwitchIgmpPortForcedLeaveState.1.2 = 2
hpSwitchIgmpPortForcedLeaveState.1.3 = 2
hpSwitchIgmpPortForcedLeaveState.1.4 = 2
hpSwitchIgmpPortForcedLeaveState.1.27 = 2
hpSwitchIgmpPortForcedLeaveState.1.28 = 2
hpSwitchIgmpPortForcedLeaveState.1.29 = 2
hpSwitchIgmpPortForcedLeaveState.1.30 = 2
hpSwitchIgmpPortForcedLeaveState.1.31 = 2
hpSwitchIgmpPortForcedLeaveState.1.32 = 2
hpSwitchIgmpPortForcedLeaveState.1.33 = 2
hpSwitchIgmpPortForcedLeaveState.1.34 = 2
hpSwitchIgmpPortForcedLeaveState.1.35 = 2
hpSwitchIgmpPortForcedLeaveState.1.36 = 2
hpSwitchIgmpPortForcedLeaveState.1.37 = 2
hpSwitchIgmpPortForcedLeaveState.1.38 = 2
hpSwitchIgmpPortForcedLeaveState.1.39 = 2
hpSwitchIgmpPortForcedLeaveState.1.40 = 2
hpSwitchIgmpPortForcedLeaveState.1.41 = 2
hpSwitchIgmpPortForcedLeaveState.1.42 = 2
hpSwitchIgmpPortForcedLeaveState.1.43 = 2
hpSwitchIgmpPortForcedLeaveState.1.44 = 2
hpSwitchIgmpPortForcedLeaveState.1.45 = 2
hpSwitchIgmpPortForcedLeaveState.1.46 = 2
hpSwitchIgmpPortForcedLeaveState.1.47 = 2
hpSwitchIgmpPortForcedLeaveState.1.48 = 2
hpSwitchIgmpPortForcedLeaveState.1.49 = 2
hpSwitchIgmpPortForcedLeaveState.1.50 = 2
```

In this example, the **2** at the end of each port listing shows that Fast Forced-Leave is disabled on all ports in the switch.

Note: In this example, the switch has a 4-port module in slot “A” and a 24-port module in slot “B”.

Figure 12-7. Listing the Forced Fast-Leave State for Ports in a Series 5300XL Switch

To list the Forced Fast-Leave state for a single port.

Syntax: getmib hpSwitchIgmpPortForcedLeaveState.1. <*port-number*>
(Not case-sensitive.)
getmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5.1. <*port-number*>

For example, to use either of the above command versions to list the state for port B1 (position 27, as shown in figure 12-7, above):

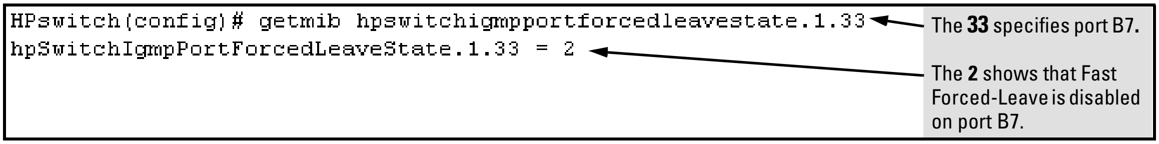

```
HPswitch(config)# getmib hpswitchigmpportforcedleavestate.1.33
hpSwitchIgmpPortForcedLeaveState.1.33 = 2
```

The **33** specifies port B7.

The **2** shows that Fast Forced-Leave is disabled on port B7.

Figure 12-8. Listing the Forced Fast-Leave State for a Single Port

Configuring Per-Port Forced Fast-Leave IGMP. In the factory-default configuration, Forced Fast-Leave is disabled for all ports on the switch. To enable (or disable) this feature on individual ports, use the switch's MIB commands, as shown below.

Syntax: setmib hpSwitchIgmpPortForcedLeaveState.1.<*port-number*> -i < 1 | 2 >
or

setmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5.1.<*port-number*> -i < 1 | 2 >

where:

1 = Forced Fast-Leave enabled

2 = Forced Fast-Leave disabled

For example, to enable Forced Fast-Leave on ports B7 and B8:

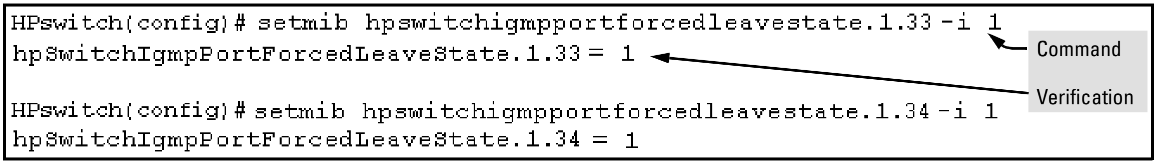

```
HPswitch(config)# setmib hpswitchigmpportforcedleavestate.1.33 -i 1
hpSwitchIgmpPortForcedLeaveState.1.33 = 1

HPswitch(config)# setmib hpswitchigmpportforcedleavestate.1.34 -i 1
hpSwitchIgmpPortForcedLeaveState.1.34 = 1
```

Command

Verification

Figure 12-9. Example of Changing the Forced Fast-Leave Configuration on Ports 7 and 8

Querier Operation

The function of the IGMP Querier is to poll other IGMP-enabled devices in an IGMP-enabled VLAN to elicit group membership information. The switch performs this function if there is no other device in the VLAN, such as a multicast router, to act as Querier. Although the switch automatically ceases Querier operation in an IGMP-enabled VLAN if it detects another Querier on the VLAN, you can also use a CLI command to disable the Querier function for that VLAN. For example, to disable the Querier function on VLAN 1 in the switch:

1. Disable Querier function on VLAN 1.

```
HPswitch(config)# no vlan 1 ip igmp querier
```

2. Disable Querier function on VLAN 1 from within the VLAN 1 context.

```
      HPswitch(vlan-1)# no ip igmp querier
```

Note

A Querier is required for proper IGMP operation. For this reason, if you disable the Querier function on the switch, ensure that there is an IGMP Querier (and, preferably, a backup Querier) available on the same VLAN.

If the switch becomes the Querier for a particular VLAN (for example, the DEFAULT_VLAN), then subsequently detects queries transmitted from another device on the same VLAN, the switch ceases to operate as the Querier for that VLAN. If this occurs, the switch Event Log lists a pair of messages similar to these:

```
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: Other Querier
detected
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: This switch is no
longer Querier
```

In the above scenario, if the other device ceases to operate as a Querier on the default VLAN, then the switch detects this change and can become the Querier as long as it is not preempted by some other IGMP Querier on the VLAN. In this case, the switch Event Log lists messages similar to the following to indicate that the switch has become the Querier on the VLAN:

```
I 01/15/01 09:21:55 igmp: DEFAULT_VLAN: Querier Election
in process
I 01/15/01 09:22:00 igmp: DEFAULT_VLAN: This switch has
been elected as Querier
```

Excluding Well-Known or Reserved Multicast Addresses from IP Multicast Filtering

Each multicast host group is identified by a single IP address in the range of 224.0.0.0 through 239.255.255.255 (which corresponds to the Ethernet multicast address range of 01005e-000000 through 01005e-7fffff). Specific groups of consecutive addresses in this range are termed "well-known" addresses and are reserved for predefined host groups. IGMP does not filter these addresses, so any packets the switch receives for such addresses are flooded out all ports assigned to the VLAN on which they were received (except the port on which the packets entered the VLAN). The following table lists the 32 well-known address groups (8192 total addresses) that IGMP does not filter on the Series 5300XL switches.

Table 12-1.Well-Known IP Multicast Address Groups Excluded from IGMP Filtering

Groups of Consecutive Addresses in the Range of 224.0.0.*X* to 239.0.0.*X**		Groups of Consecutive Addresses in the Range of 224.128.0.*X* to 239.128.0.*X**	
224.0.0.*x*	232.0.0.*x*	224.128.0.*x*	232.128.0.*x*
225.0.0.*x*	233.0.0.*x*	225.128.0.*x*	233.128.0.*x*
226.0.0.*x*	234.0.0.*x*	226.128.0.*x*	234.128.0.*x*
227.0.0.*x*	235.0.0.*x*	227.128.0.*x*	235.128.0.*x*
228.0.0.*x*	236.0.0.*x*	228.128.0.*x*	236.128.0.*x*
229.0.0.*x*	237.0.0.*x*	229.128.0.*x*	237.128.0.*x*
230.0.0.*x*	238.0.0.*x*	230.128.0.*x*	238.128.0.*x*
231.0.0.*x*	239.0.0.*x*	231.128.0.*x*	239.128.0.*x*

* *X* is any value from 0 to 255.

Devices such as the HP Procurve Series 5300XL and the HP Procurve switch 1600M/2400M/2424M/4000M/8000M having static Traffic/Security filters configured with a "Multicast" filter type and a "Multicast Address" in this range will continue in effect unless IGMP learns of a multicast group destination in this range. In that case, IGMP takes over the filtering function for the multicast destination address(es) for as long as the IGMP group is active. If the IGMP group subsequently deactivates, the static filter resumes control over traffic to the multicast address formerly controlled by IGMP.

802.1w Rapid Spanning Tree Protocol (RSTP) 802.1d Spanning Tree Protocol (STP)

Contents

Overview

STP Features

802.1d Spanning Tree Protocol	**Default**	**Menu**	**CLI**	**Web**
viewing the STP configuration	n/a	page 13-19	page 13-10	—
enable/disable STP	disabled	page 13-19	page 13-23	page 13-41
reconfiguring general operation	priority: 32768 max age: 20 s hello time: 2 s fwd. delay: 15 s	page 13-19	page 13-24	
reconfiguring per-port STP	path cost: var priority: 128 mode: norm	page 13-19	page 13-25	
monitoring STP	n/a	page B-16	page B-16	n/a
802.1w Spanning Tree Protocol	**Default**	**Menu**	**CLI**	**Web**
Viewing the RSTP/STP configuration	--	page 13-16	page 13-10	n/a
enable/disable RSTP/STP (RSTP is selected as the default protocol)	disabled	page 13-16	page 13-11	page 13-18
reconfiguring whole-switch values	Protocol Version: RSTP Force Version: RSTP-operation Switch Priority: 8 Hello Time: 2 s Max Age: 20 s Forward Delay: 15 s	page 13-16	page 13-12	n/a
reconfiguring per-port values	Path Cost: Depends on port type Priority: 8 Edge Port: Yes Point-to-point: Force-true MCheck: Yes	page 13-16	page 13-14	n/a

Use spanning tree to ensure that only one active path at a time exists between any two nodes on the network. In networks where there is more than one physical, active path between any two nodes, enabling spanning tree ensures a single active path between such nodes by blocking all redundant paths. Without spanning tree, having more than one active path between a pair of nodes causes loops in the network, which can result in duplication of messages, leading to a "broadcast storm" that can bring down the network.

Note

You should enable spanning tree operation in any switch that is part of a redundant physical link (loop topology). (It is recommended that you do so on all switches belonging to a loop topology.) This topic is covered in more detail under "How Spanning Tree Operates" on page 13-4.

As recommended in the IEEE 802.1Q VLAN standard, the Series 5300XL switches use **single-instance STP**. (As a result, the switch generates untagged Bridge Protocol Data Units—BPDUs.) This implementation creates a single spanning tree to make sure there are no network loops associated with any of the connections to the switch, regardless of whether multiple VLANs are configured on the switch. Thus, these switches do not distinguish between VLANs when identifying redundant physical links. If VLANs are configured on the switch, see "STP Operation with 802.1Q VLANs" on page "Spanning Tree Operation with 802.1Q VLANs" on page 13-4.

How Spanning Tree Operates

The switch automatically senses port identity and type, and automatically defines spanning-tree parameters for each type, as well as parameters that apply across the switch. You can use the default values for these parameters, or adjust them as needed.

While allowing only one active path through a network at any time, spanning tree retains any redundant physical path to serve as a backup (blocked) path in case the existing active path fails. Thus, if an active path fails, spanning tree automatically activates (unblocks) an available backup to serve as the new active path for as long as the original active path is down. For example:

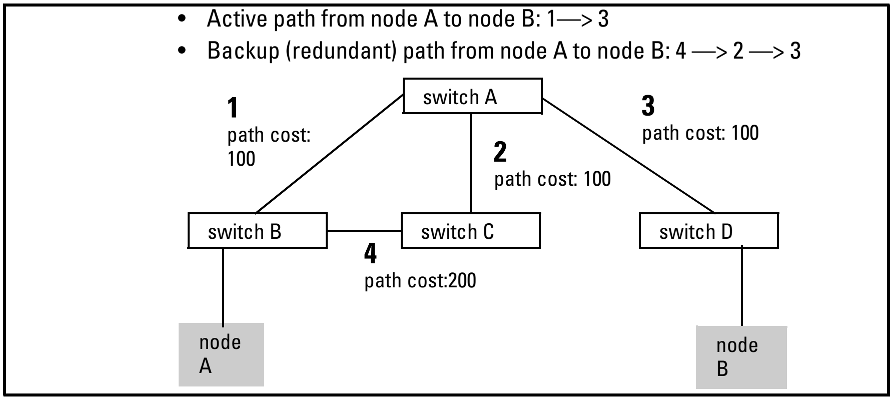

- Active path from node A to node B: 1—> 3
- Backup (redundant) path from node A to node B: 4 —> 2 —> 3

Figure 13-1. General Example of Redundant Paths Between Two Nodes

In the factory default configuration, spanning tree operation is off. If a redundant link (loop) exists between nodes in your network, you should enable the spanning tree operation of your choice.

Note

Spanning tree retains its current parameter settings when disabled. Thus, if you disable spanning tree, then later re-enable it, the parameter settings will be the same as before spanning tree was disabled.

Spanning Tree Operation with 802.1Q VLANs. As recommended in the IEEE 802.1Q VLAN standard, when spanning tree is enabled on the switch, a single spanning tree is configured for all ports across the switch, including those in separate VLANs (that is, single-instance spanning tree, which generates untagged BPDUs). This means that if redundant physical links exist in separate VLANs, spanning tree will block all but one of those links. However, if you need to use spanning tree on the Series 5300XL switches in a VLAN environment with redundant physical links, you can prevent blocked redundant links by using a port trunk. The following example shows how you can use a port trunk with 802.1Q (tagged) VLANs and spanning tree without unnecessarily blocking any links or losing any bandwidth.

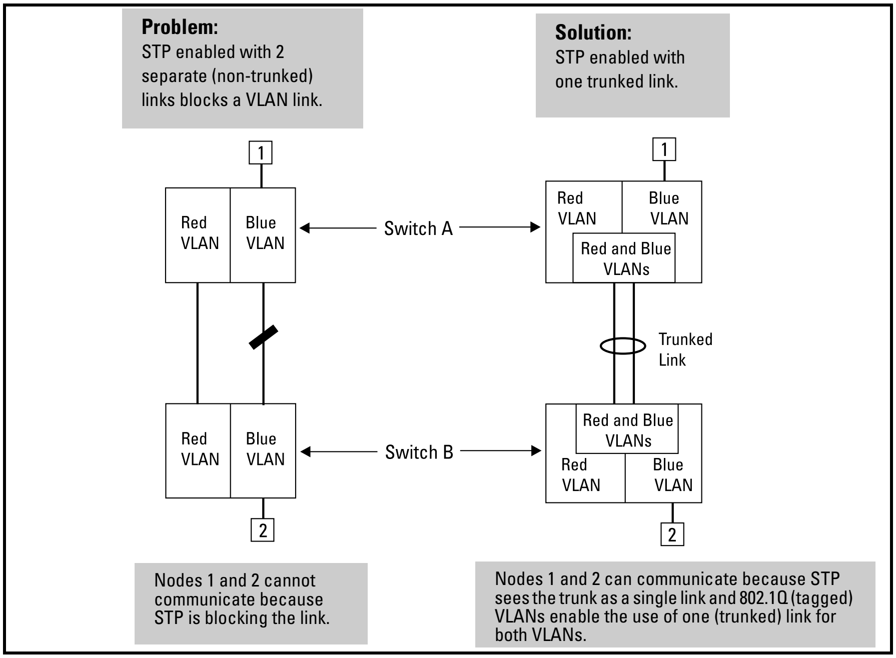
Problem:
STP enabled with 2 separate (non-trunked) links blocks a VLAN link.
Solution:
STP enabled with one trunked link.
1
Red VLAN
Blue VLAN
Switch A
Red VLAN
Blue VLAN
Red and Blue VLANs
Trunked Link
Red and Blue VLANs
Red VLAN
Blue VLAN
Switch B
2
Nodes 1 and 2 cannot communicate because STP is blocking the link.
Nodes 1 and 2 can communicate because STP sees the trunk as a single link and 802.1Q (tagged) VLANs enable the use of one (trunked) link for both VLANs.

Figure 13-2. Example of Using a Trunked Link with STP and VLANs

For more information, refer to “Spanning Tree Operation with 802.1Q VLANs” on page 13-4.

Spanning Tree Options: RSTP (802.1w) and STP (802.1d)

RSTP (802.1w)

The IEEE 802.1d version of spanning tree (STP) can take a fairly long time to resolve all the possible paths and to select the most efficient path through the network. The IEEE 802.1w Rapid Reconfiguration Spanning Tree (RSTP) significantly reduces the amount of time it takes to establish the network path. The result is reduced network downtime and improved network robustness.

In addition to faster network reconfiguration, RSTP also implements greater ranges for port path costs to accommodate the higher and higher connection speeds that are being implemented.

RSTP is designed to be compatible with IEEE 802.1d STP, and HP recommends that you employ it in your network. For more information, refer to "Transitioning from STP to RSTP" on page 13-8.

STP (802.1d)

The IEEE 802.1d version of spanning tree has been in wide use and can coexist in a network in which RSTP (802.1w) has been introduced. if your network currently uses 802.1d STP and you are not yet ready to implement RSTP, you can apply STP to the Series 5300XL switches until such time as you are ready to move ahead with RSTP. STP on the Series 5300XL switches offers the full range of STP features found in earlier product releases, including:

- **STP Fast Mode for Overcoming Server Access Failures:** If an end node is configured to automatically access a server, the duration of the STP startup sequence can result in a "server access failure". On ports where this is a problem, configuring STP Fast Mode can eliminate the failure. For more information, see "STP Fast Mode" on page 13-26. The next sections describe how to configure STP on the switch. For more information on STP operation, see "How Spanning Tree Operates" on page 13-4.
- **Fast-Uplink STP for Improving the Recovery (Convergence) Time in Wiring Closet Switches with Redundant Uplinks:** This means that a Series 5300XL switch having redundant links toward the root device can decrease the convergence time to a new uplink port to as little as ten seconds. For more information, refer to "Fast-Uplink Spanning Tree Protocol (STP)" on page 13-27.

Caution

Because the switch automatically gives faster links a higher priority, the default STP parameter settings are usually adequate for spanning tree operation. Also because incorrect STP settings can adversely affect network performance, you should not make changes unless you have a strong understanding of how STP operates. For more on STP, see the IEEE 802.1D standard.

Configuring Rapid Reconfiguration Spanning Tree (RSTP)

This section describes the operation of the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP)

Overview

RSTP Feature	Default		Menu	CLI	Web
Viewing the RSTP/STP configuration	*n/a*		page 13-16	page 13-10	n/a
enable/disable RSTP/STP (RSTP is selected as the default protocol)	disabled		page 13-16	page 13-11	page 13-18
reconfiguring whole-switch values	Protocol Version:	RSTP	page 13-16	page 13-12	n/a
	Force Version:	RSTP-operation			
	Switch Priority:	8			
	Hello Time:	2 s			
	Max Age:	20 s			
	Forward Delay:	15 s			
reconfiguring per-port values	Path Cost:	*depends on port type*	page 13-16	page 13-14	n/a
	Priority:	8			
	Edge Port:	Yes			
	Point-to-point:	Force-true			
	MCheck:	Yes			

As indicated in the manual, the spanning tree protocol is used to ensure that only one active path at a time exists between any two end nodes in the network in which your switch is installed. Multiple paths cause a loop in the network over which broadcast and multicast messages are repeated continuously, which floods the network with traffic creating a broadcast storm.

In networks where there is more than one physical path between any two nodes, enabling spanning tree ensures a single active path between two such nodes by selecting the one most efficient path and blocking the other redundant paths. If a switch or bridge in the path becomes disables, spanning tree activates the necessary blocked segments to create the next most efficient path.

Transitioning from STP to RSTP

IEEE 802.1w RSTP is designed to be compatible with IEEE 802.1d STP. Even if all the other devices in your network are using STP, you can enable RSTP on your switch, and even using the default configuration values, your switch will interoperate effectively with the STP devices. If any of the switch ports are connected to switches or bridges on your network that do not support RSTP, RSTP can still be used on this switch. RSTP automatically detects when the switch ports are connected to non-RSTP devices in the spanning tree and communicates with those devices using 802.1d STP BPDU packets.

Because RSTP is so much more efficient at establishing the network path, though, that it is highly recommended that all your network devices be updated to support RSTP. RSTP offers convergence times of less than one second under optimal circumstances. To make the best use of RSTP and achieve the fastest possible convergence times, though, there are some changes that you should make to the RSTP default configuration. See "Optimizing the RSTP Configuration" below, for more information on these changes.

Note

Under some circumstances, it is possible for the rapid state transitions employed by RSTP to result in an increase in the rates of frame duplication and misordering in the switched LAN. In order to allow RSTP switches to support applications and protocols that may be sensitive to frame duplication and misordering, setting the Force Protocol Version parameter to **STP-compatible** allows RSTP to be operated with the rapid transitions disabled. The value of this parameter applies to all ports on the switch. See information on Force Version on page 13-12.

As indicated above, one of the benefits of RSTP is the implementation of a larger range of port path costs, which accommodates higher network speeds. New default values have also been implemented for the path costs associated with the different network speeds. This can create some incompatibility between devices running the older 802.1d STP and your switch running RSTP. Please see the "Note on Path Cost" on page 13-15 for more information on adjusting to this incompatibility.

Configuring RSTP

The default switch configuration has spanning tree disabled with RSTP as the selected protocol. That is, when spanning tree is enabled, RSTP is the version of spanning tree that is enabled, by default.

Optimizing the RSTP Configuration

To optimize the RSTP configuration on your switch, follow these steps (note that for the **Menu** method, all of these steps can be performed at the same time by making all the necessary edits on the "Spanning Tree Operation" screen and then saving the configuration changes):

1. Set the switch to support RSTP (RSTP is the default):

   CLI: spanning-tree protocol-version rstp

   Menu: Main Menu —> 2. Switch Configuration —> 4. Spanning Tree Operation —> select Protocol Version: RSTP

2. Set the "point-to-point-mac" value to false on all ports that are connected to shared LAN segments (that is, to connections to hubs):

   CLI: spanning-tree [ethernet] <*port-list*> point-to-point-mac force-false

   Menu: Main Menu —> 2. Switch Configuration —> 4. Spanning Tree Operation —> for each appropriate port, select Point-to-Point: Force-False

3. Set the "edge-port" value to false for all ports connected to other switches, bridges, and hubs:

   CLI: no spanning-tree [ethernet] <*port-list*> edge-port

   Menu: Main Menu —> 2. Switch Configuration —> 4. Spanning Tree Operation —> for each appropriate port, select Edge: No

4. Set the "mcheck" value to false for all ports that are connected to devices that are known to be running IEEE 802.1d spanning tree:

   CLI: no spanning-tree [ethernet] <*port-list*> mcheck

   Menu: Main Menu —> 2. Switch Configuration —> 4. Spanning Tree Operation —> for each appropriate port, select MCheck: No

5. Enable RSTP Spanning Tree:

   CLI: spanning-tree

   Menu: Main Menu —> 2. Switch Configuration —> 4. Spanning Tree Operation —> select
   STP Enabled: Yes

CLI: Configuring RSTP

Spanning Tree Commands in This Section	**Applicable Protocol Version**	**Location**
show spanning-tree config	both	Below on this page
spanning-tree	both	page 13-11
protocol-version <rstp | stp>	both	page 13-12
force-version <rstp-operation | stp-compatible>	RSTP	page 13-12
forward-delay <4 - 30>	both	page 13-12
hello-time <1 - 10>	both	page 13-12
maximum-age <6 - 40>	both	page 13-12
priority <0 - 15 | 0 - 65535>	RSTP | STP	page 13-12
<[ethernet] *port-list*>	both	page 13-14
path-cost <1 - 200 000 000>	both	page 13-14
priority <0 - 15 | 0 - 65535>	RSTP | STP	page 13-14
edge-port	RSTP	page 13-14
point-to-point-mac	RSTP	page 13-14
mcheck	RSTP	page 13-14
mode <norm | fast>	STP	Refer to "802.1p Spanning-Tree Protocol (STP)" on page 13-19.
show spanning-tree	This command lists additional RSTP/STP monitoring data that is not covered in this section. See "Spanning Tree Protocol Information" on page B-16	

Viewing the Current Spanning Tree Configuration. Even if spanning tree is disabled (the default configuration), the show spanning-tree config command lists the switch's full spanning tree configuration, including whole-switch and per-port settings.

Syntax: show spanning-tree configuration

Abbreviation: sho span config

In the default configuration, the output from this command appears similar to the following:

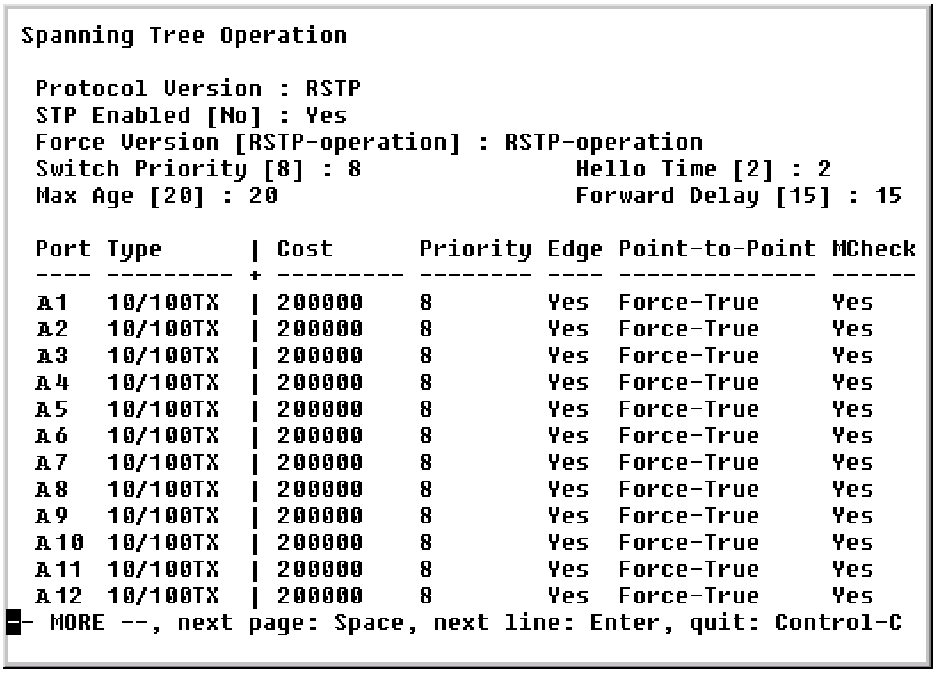

```
Spanning Tree Operation

 Protocol Version : RSTP
 STP Enabled [No] : Yes
 Force Version [RSTP-operation] : RSTP-operation
 Switch Priority [8] : 8               Hello Time [2] : 2
 Max Age [20] : 20                     Forward Delay [15] : 15

 Port Type      | Cost      Priority Edge Point-to-Point MCheck
 ---- --------- + --------- -------- ---- -------------- ------
 A1   10/100TX  | 200000    8        Yes  Force-True     Yes
 A2   10/100TX  | 200000    8        Yes  Force-True     Yes
 A3   10/100TX  | 200000    8        Yes  Force-True     Yes
 A4   10/100TX  | 200000    8        Yes  Force-True     Yes
 A5   10/100TX  | 200000    8        Yes  Force-True     Yes
 A6   10/100TX  | 200000    8        Yes  Force-True     Yes
 A7   10/100TX  | 200000    8        Yes  Force-True     Yes
 A8   10/100TX  | 200000    8        Yes  Force-True     Yes
 A9   10/100TX  | 200000    8        Yes  Force-True     Yes
 A10  10/100TX  | 200000    8        Yes  Force-True     Yes
 A11  10/100TX  | 200000    8        Yes  Force-True     Yes
 A12  10/100TX  | 200000    8        Yes  Force-True     Yes
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

Figure 13-3. Example of the Spanning Tree Configuration Display

Enabling or Disabling RSTP. Issuing the command to enable spanning tree on the switch implements, by default, the RSTP version of spanning tree for all physical ports on the switch. Disabling spanning tree removes protection against redundant network paths.

Syntax: [no] spanning-tree

> ***Abbreviation:*** [no] span

This command enables spanning tree with the current parameter settings or disables spanning tree, using the “no” option, without losing the most-recently configured parameter settings.

Enabling STP Instead of RSTP. If you decide, for whatever reason, that you would prefer to run the IEEE 802.1d (STP) version of spanning tree, then issue the following command:

Syntax: spanning-tree protocol-version stp

> ***Abbreviation:*** span prot stp

For the STP version of spanning tree, the rest of the information in this section does not apply. Refer to "802.1p Spanning-Tree Protocol (STP)" on page 13-19 for more information on the STP version and its parameters.

Reconfiguring Whole-Switch Spanning Tree Values. You can configure one or more of the following parameters, which affect the spanning tree operation of the whole switch:

Table 13-1.Whole-Switch RSTP Parameters

Parameter	Default	Description
protocol-version	RSTP	Identifies which of the spanning tree protocols will be used when spanning tree is enabled on the switch.
force-version	rstp-operation	Sets the spanning tree compatibility mode. Even if **rstp-operation** is selected though, if the switch detects STP BPDU packets on a port, it will communicate to the attached device using STP BPDU packets. If errors are encountered, as described in the Note on page 8, the Force-Version value can be set to **stp-compatible**, which forces the switch to communicate out all ports using operations that are compatible with IEEE 802.1d STP.
priority	32768 (8 as a step value)	Specifies the protocol value used along with the switch MAC address to determine which device in the spanning tree is the root. The lower the priority value, the higher the priority. The value you enter has changed from the STP value. The range is 0 - 61440, but for RSTP the value is entered as a multiple (a step) of 4096. You enter a value in the range 0 - 15. The default value of 32768 is derived by the default setting of 8. Displaying the RSTP configuration (**show spanning-tree config**) shows 8, but displaying the RSTP operation (**show spanning-tree**) shows 32768.
*maximum-age	20 seconds	Sets the maximum age of received spanning tree information before it is discarded. The range is 6 to 40 seconds.
*hello-time	2 seconds	Sets the time between transmission of spanning tree messages. Used only when this switch is the root. The range is 1 to 10 seconds.
*forward-delay	15 seconds	Sets the time the switch waits between transitioning ports from listening to learning and from learning to forwarding states. The range is 4 to 30 seconds.

*These parameters are the same for RSTP as they are for STP. The switch uses its own maximum-age, hello-time, and forward-delay settings only if it is operating as the root device in the spanning tree. If another device is the root device, then the switch uses the other device's settings for these parameters.

Note

Executing the **spanning-tree** command alone enables spanning tree. Executing the command with one or more of the whole-switch RSTP parameters shown in the table on the previous page, or with any of the per-port RSTP parameters shown in the table on page 14, does not enable spanning tree. It only configures the spanning tree parameters, regardless of whether spanning tree is actually running (enabled) on the switch.

Using this facility, you can completely configure spanning tree the way you want and then enable it. This method minimizes the impact on the network operation.

Syntax:	***Abbreviations:***
spanning-tree	span
protocol-version <rstp | stp>	prot <rstp | stp>
force-version <rstp-operation | stp-compatible>	forc <rstp | stp>
priority <0 - 15>	pri <0 - 15>
maximum-age <6 - 40 seconds>	max <6 - 40>
hello-time <1- 10 seconds>	hello <1 - 10>
forward-delay <4 - 30 seconds>	forw <4 - 30>

Defaults: see the table on the previous page.

Multiple parameters can be included on the same command line. For example, to configure a maximum-age of 30 seconds and a hello-time of 3 seconds, you would issue the following command:

```
HPswitch (config)# span max 30 hello 3k
```

Reconfiguring Per-Port Spanning Tree Values. You can configure one or more of the following parameters, which affect the spanning tree operation of the specified ports only:

Table 13-2.Per-Port RSTP Parameters

Parameter	Default	Description
edge-port	Yes	Identifies ports that are connected to end nodes. During spanning tree establishment, these ports transition immediately to the Forwarding state. In this way, the ports operate very similarly to ports that are configured in "fast mode" under the STP implementation in previous HP switch software. Disable this feature on all switch ports that are connected to another switch, or bridge, or hub. Use the "no" option on the spanning tree command to disable edge-port.
mcheckt	Yes	Ports with mcheck set to true are forced to send out RSTP BPDUs for 3 seconds. This allows for switches that are running RSTP to establish their connection quickly and for switches running 802.1d STP to be identified. If the whole-switch parameter Force-Version is set to "stp-compatible", the mcheck setting is ignored and STP BPDUs are sent out all ports. Disable this feature on all ports that are known to be connected to devices that are running 802.1d STP. Use the "no" option on the spanning tree command to disable mcheck.
path-cost	10 Mbps – 2 000 000 100 Mbps – 200 000 1 Gbps – 20 000	Assigns an individual port cost that the switch uses to determine which ports are the forwarding ports. The range is 1 to 200,000,000 or auto. By default, this parameter is automatically determined by the port type, as shown by the different default values. If you have previously configured a specific value for this parameter, you can issue the command with the **auto** option to restore the automatic setting feature. Please see the Note on Path Cost on page 13-15 for information on compatibility with devices running 802.1d STP for the path cost values.
point-to-point-mac	force-true	This parameter is used to tell the port if it is connected to a point-to-point link, such as to another switch or bridge or to an end node (**force-true**). This parameter should be set to **force-false** for all ports that are connected to a hub, which is a shared LAN segment. You can also set this parameter to **auto** and the switch will automatically set the force-false value on all ports that it detects are not running at full duplex. All connections to hubs are not full duplex.
priority	128 (8 as a step value)	This parameter is used by RSTP to determine the port(s) to use for forwarding. The port with the lowest number has the highest priority. The range is 0 to 240, but you configure the value by entering a multiple of 16. You enter a value in the range 0 - 15. The default value of 128 is derived by the default setting of 8. Displaying the RSTP configuration (**show spanning-tree config**) shows 8, but displaying the RSTP operation (**show spanning-tree**) shows 128.

Syntax:	***Abbreviations:***
spanning-tree [ethernet] <*port-list*>	span <*port-list*>
path-cost <1 - 200000000>	path <1 - 200000000>
point-to-point-mac <force-true | force-false | auto>	forc <force-t | force-f | auto>
priority <0 - 15>	pri <0 - 15>
[no] spanning-tree [ethernet] <*port-list*>	[no] span <port-list>
edge-port	edge
mcheck	mch

Defaults: see the table on the previous page.

Note on Path Cost

RSTP implements a greater range of path costs and new default path cost values to account for higher network speeds. These values are different than the values defined by 802.1d STP as shown below.

Port Type	802.1d STP Path Cost	RSTP Path Cost
10 Mbps	100	2 000 000
100 Mbps	10	200 000
1 Gbps	5	20 000
10 Gbps	?	2000

Because the maximum value for the path cost allowed by 802.1d STP is 65535, devices running that version of spanning tree cannot be configured to match the values defined by RSTP, at least for 10 Mbps and 100 Mbps ports. In LANs where there is a mix of devices running 802.1d STP and RSTP, you should reconfigure the devices so the path costs match for ports with the same network speeds.

Menu: Configuring RSTP

1. From the console CLI prompt, enter the menu command.

   HP Procurve Switch # **menu**

2. From the switch console Main Menu, select

   2. Switch Configuration ...

   4. Spanning Tree Operation

3. Press [E] (for **Edit**) to highlight the **Protocol Version** parameter field.

4. Press the Space bar to select the version of spanning tree you wish to run: **RSTP** or **STP**.

   Note: If you change the protocol version, you will have to reboot the switch for the change to take effect. See step 9 and step 10.

5. Press the [Tab] or down arrow key to go to the **STP Enabled** field. Note that when you do this, the remaining fields on the screen will then be appropriate for the version of spanning tree that was selected in step 3. The screen image below is for RSTP.

6. Press the Space bar to select **Yes** to enable spanning tree.

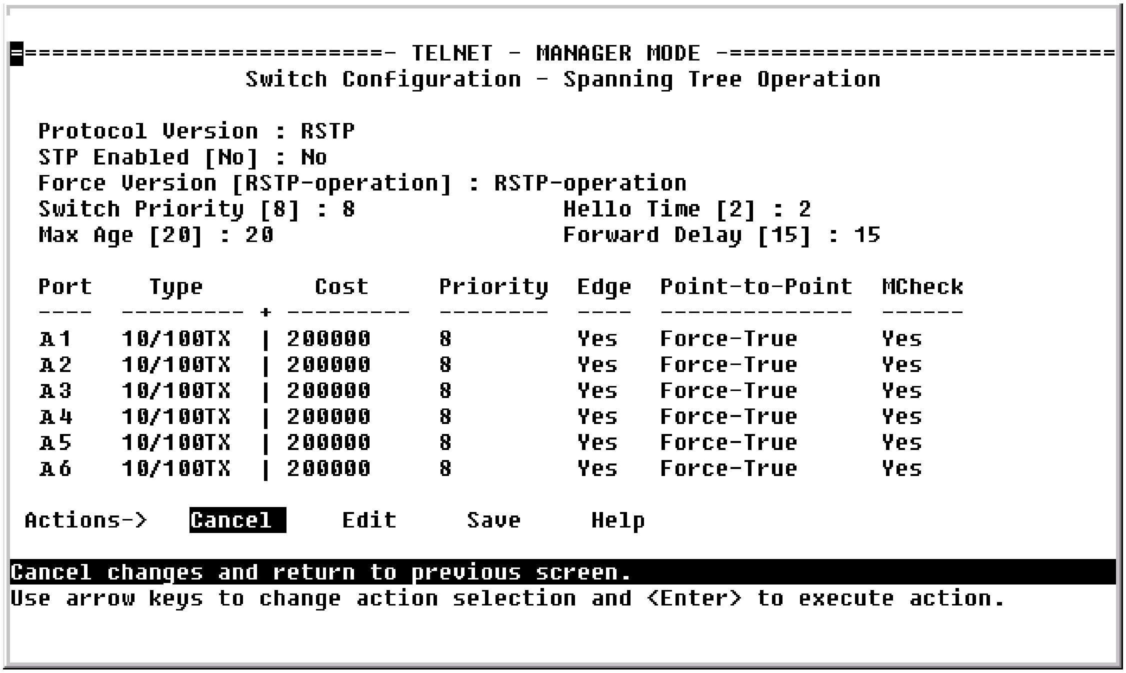

```
=========================- TELNET - MANAGER MODE -============================
                Switch Configuration - Spanning Tree Operation

  Protocol Version : RSTP
  STP Enabled [No] : No
  Force Version [RSTP-operation] : RSTP-operation
  Switch Priority [8] : 8              Hello Time [2] : 2
  Max Age [20] : 20                    Forward Delay [15] : 15

  Port    Type        Cost      Priority  Edge  Point-to-Point  MCheck
  ----  --------- + ---------  --------  ----  --------------  ------
  A1    10/100TX  | 200000     8         Yes   Force-True      Yes
  A2    10/100TX  | 200000     8         Yes   Force-True      Yes
  A3    10/100TX  | 200000     8         Yes   Force-True      Yes
  A4    10/100TX  | 200000     8         Yes   Force-True      Yes
  A5    10/100TX  | 200000     8         Yes   Force-True      Yes
  A6    10/100TX  | 200000     8         Yes   Force-True      Yes

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 13-4. Example of the RSTP Configuration Screen

7. Press the [Tab] key or use the arrow keys to go to the next parameter you want to change, then type in the new value or press the Space bar to select a value. (To get help on this screen, press [Enter] to select the **Actions –>** line, then press [H], for **Help**, to display the online help.)

8. Repeat step 6 for each additional parameter you want to change.

   Please see “Optimizing the RSTP Configuration” on page 13-9 for recommendations on configuring RSTP to make it operate the most efficiently.

9. When you are finished editing parameters, press [Enter] to return to the **Actions –>** line and press [S] to save the currently displayed spanning tree settings and return to the Main Menu.

10. If you have changed the Protocol Version, in step 1, reboot the switch now by selecting

    6. Reboot Switch

Web: Enabling or Disabling RSTP

In the web browser interface, you can enable or disable spanning tree on the switch. If the default configuration is in effect such that RSTP is the selected protocol version, enabling spanning tree through the web browser interface will enable RSTP with its current configuration. To configure the other spanning tree features, telnet to the switch console and use the CLI or menu.

To enable or disable spanning tree using the web browser interface:

1. Click on the **Configuration** tab.
2. Click on [Device Features].
3. Enable or disable spanning tree.
4. Click on [Apply Changes] to implement the configuration change.

802.1p Spanning-Tree Protocol (STP)

Menu: Configuring 802.1D STP

1. From the Main Menu, select:

   2. Switch Configuration . . .

   4. Spanning Tree Operation

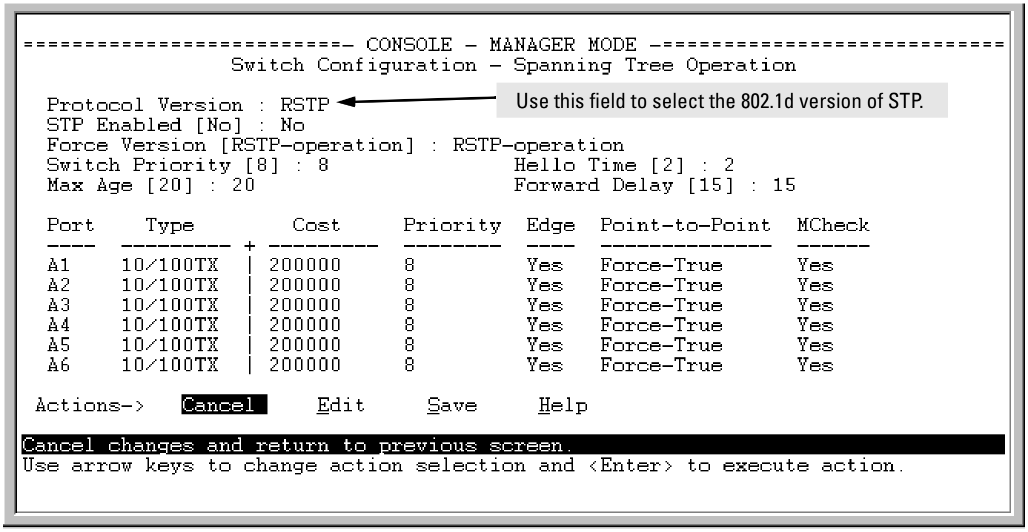

```
========================- CONSOLE - MANAGER MODE -============================
                 Switch Configuration - Spanning Tree Operation

  Protocol Version : RSTP
  STP Enabled [No] : No
  Force Version [RSTP-operation] : RSTP-operation
  Switch Priority [8] : 8                Hello Time [2] : 2
  Max Age [20] : 20                      Forward Delay [15] : 15

  Port    Type        Cost      Priority  Edge  Point-to-Point  MCheck
  ----  --------- + ---------  --------  ----  --------------  ------
  A1    10/100TX  | 200000     8         Yes   Force-True      Yes
  A2    10/100TX  | 200000     8         Yes   Force-True      Yes
  A3    10/100TX  | 200000     8         Yes   Force-True      Yes
  A4    10/100TX  | 200000     8         Yes   Force-True      Yes
  A5    10/100TX  | 200000     8         Yes   Force-True      Yes
  A6    10/100TX  | 200000     8         Yes   Force-True      Yes

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Use this field to select the 802.1d version of STP.

Figure 13-5. The Default "Spanning Tree Operation" Screen

2. Press [E] (for **Edit**) to highlight the **Protocol Version** field. In the default configuration this field is set to **RSTP**.

3. Press the Space bar once to change the field to STP. This changes the Protocol Version selection to the 802.1d Spanning Tree Protocol.

4. Press [↓] to highlight the **STP Enabled** field.

5. Press the Space bar to select **Yes**. (**Yes** in this field means to enable spanning-tree operation.)

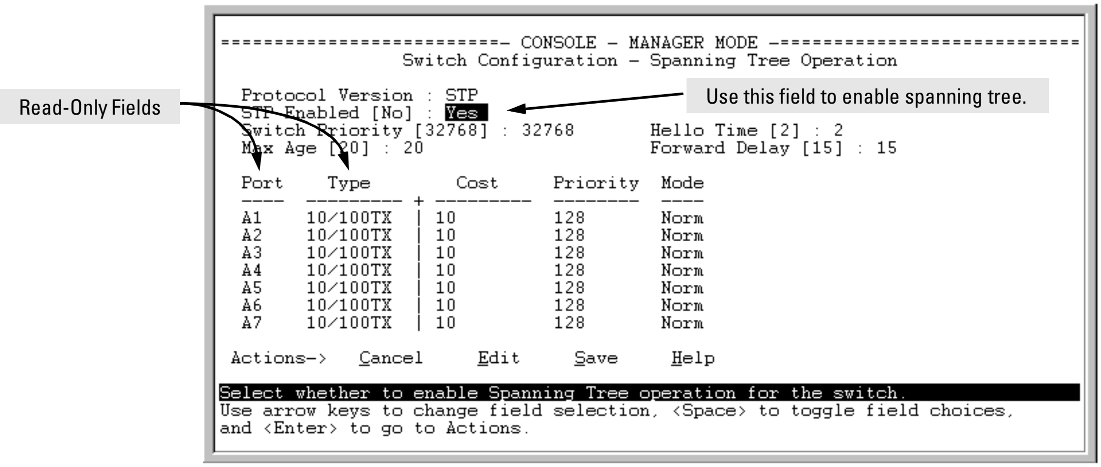

```
==========================- CONSOLE - MANAGER MODE -===========================
                 Switch Configuration - Spanning Tree Operation

  Protocol Version : STP
  STP Enabled [No] : Yes
  Switch Priority [32768] : 32768       Hello Time [2] : 2
  Max Age [20] : 20                     Forward Delay [15] : 15

  Port    Type        Cost      Priority  Mode
  ----    --------- + --------- --------  ----
  A1      10/100TX  | 10        128       Norm
  A2      10/100TX  | 10        128       Norm
  A3      10/100TX  | 10        128       Norm
  A4      10/100TX  | 10        128       Norm
  A5      10/100TX  | 10        128       Norm
  A6      10/100TX  | 10        128       Norm
  A7      10/100TX  | 10        128       Norm

 Actions->   Cancel     Edit     Save     Help

Select whether to enable Spanning Tree operation for the switch.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Read-Only Fields

Use this field to enable spanning tree.

Figure 13-6. Enabling Spanning-Tree Operation

6. If the remaining STP parameter settings are adequate for your network, go to step 10.

7. Use Tab or the arrow keys to select the next parameter you want to change, then type in the new value or press the Space Bar to select a value. (If you need information on STP parameters, press Enter to select the **Actions** line, then press **H** to get help.)

8. Repeat step 7 for each additional parameter you want to change.

   Note: For information on the **Mode** parameter, see "STP Fast Mode" on page 13-26.

9. When you are finished editing parameters, press Enter to return to the **Actions** line.

10. Press S to save the currently displayed STP parameter settings. You will then see the "Switch Configuration Menu" with an asterisk (*) at the **Spanning Tree Operation** line, indicating that you must reboot the switch before the Protocol Version change (step 5) takes effect.

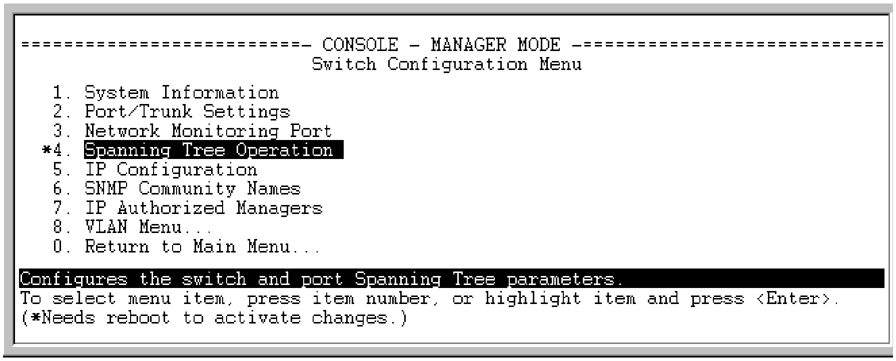

```
==========================- CONSOLE - MANAGER MODE -============================
                         Switch Configuration Menu

   1. System Information
   2. Port/Trunk Settings
   3. Network Monitoring Port
  *4. Spanning Tree Operation
   5. IP Configuration
   6. SNMP Community Names
   7. IP Authorized Managers
   8. VLAN Menu...
   0. Return to Main Menu...

Configures the switch and port Spanning Tree parameters.
To select menu item, press item number, or highlight item and press <Enter>.
(*Needs reboot to activate changes.)
```

Figure 13-7. The Configuration Menu Indicating a Reboot Is Needed to Implement a Configuration Change

11. Press [0] to return to the Main menu.

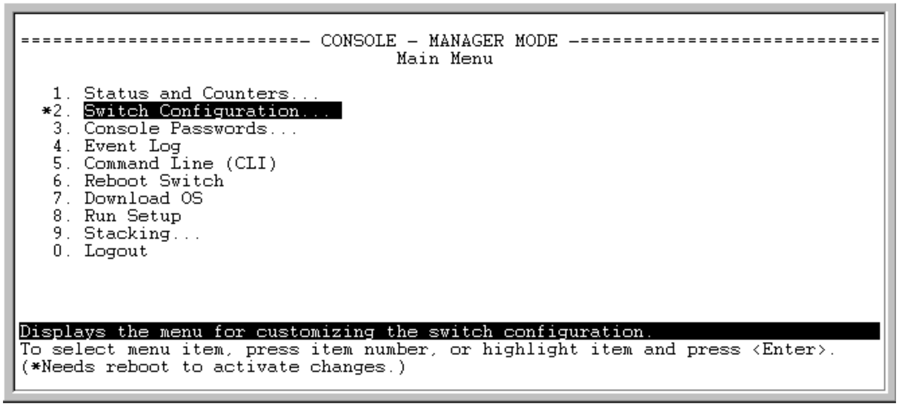

```
==========================- CONSOLE - MANAGER MODE -============================
                                Main Menu

   1. Status and Counters...
  *2. Switch Configuration...
   3. Console Passwords...
   4. Event Log
   5. Command Line (CLI)
   6. Reboot Switch
   7. Download OS
   8. Run Setup
   9. Stacking...
   0. Logout



Displays the menu for customizing the switch configuration.
To select menu item, press item number, or highlight item and press <Enter>.
(*Needs reboot to activate changes.)
```

Figure 13-8. The Main Menu Indicating a Reboot Is Needed To Implement a Configuration Change

12. Press [6] to reboot the switch. This implements the Protocol Version change (steps 2 and 3 on page 13-19).

CLI: Configuring 802.1D STP

STP Commands Used in This Section

show spanning-tree config	Below
spanning-tree	
protocol-version	page 13-23
forward-delay *<4 - 30>*	page 13-24
hello-time *<1 - 10>*	page 13-24
maximum-age *<6 - 40>*	page 13-24
priority *<0 - 65535>*	page 13-24
ethernet *<port-list>*	page 13-25
path-cost *<1 - 65535>*	page 13-25
priority *<0 - 255>*	page 13-25
mode *<norm | fast>*	page 13-25
show spanning tree	Lists additional STP data not covered in this chapter. See "Spanning Tree Protocol (STP) Information" on page B-16

Viewing the Current STP Configuration. Regardless of whether STP is disabled (the default), this command lists the switch's full STP configuration, including general settings and port settings.

Syntax: show spanning-tree config

When the switch is configured for 802.1d STP, this command displays information similar to the following:

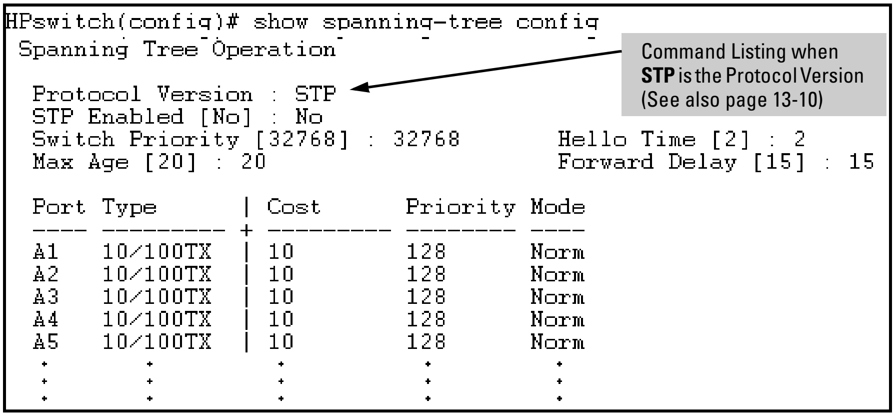

```
HPswitch(config)# show spanning-tree config
 Spanning Tree Operation

  Protocol Version : STP
  STP Enabled [No] : No
  Switch Priority [32768] : 32768       Hello Time [2] : 2
  Max Age [20] : 20                     Forward Delay [15] : 15

  Port Type      | Cost      Priority Mode
  ---- --------- + --------- -------- ----
  A1   10/100TX  | 10        128      Norm
  A2   10/100TX  | 10        128      Norm
  A3   10/100TX  | 10        128      Norm
  A4   10/100TX  | 10        128      Norm
  A5   10/100TX  | 10        128      Norm
   .      .         .          .        .
   .      .         .          .        .
   .      .         .          .        .
```

Command Listing when **STP** is the Protocol Version (See also page 13-10)

Figure 13-9. Example of the Default STP Configuration Listing with 802.1d STP Configured at the Protocol Version

Configuring the Switch To Use the 802.1d Spanning Tree Protocol (STP). In the default configuration, the switch is set to **RSTP** (that is, 802.1w Rapid Spanning Tree), and spanning tree operation is disabled. To reconfigure the switch to 802.1d spanning tree, you must:

1. Change the spanning tree protocol version to **stp**.
2. Use **write memory** to save the change to the startup-configuration.
3. Reboot the switch.
4. If you have not previously enabled spanning-tree operation on the switch, use the **spanning-tree** command again to enable STP operation.

Syntax: spanning-tree protocol-version stp
write memory
boot

For example:

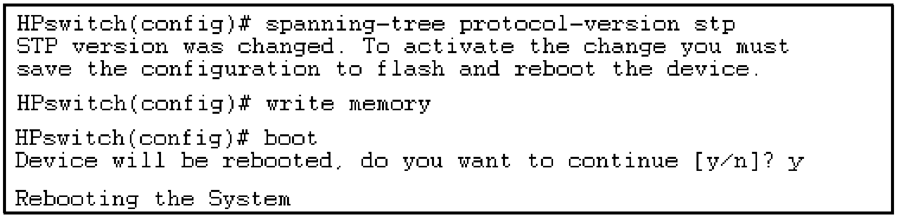

```
HPswitch(config)# spanning-tree protocol-version stp
STP version was changed. To activate the change you must
save the configuration to flash and reboot the device.

HPswitch(config)# write memory

HPswitch(config)# boot
Device will be rebooted, do you want to continue [y/n]? y

Rebooting the System
```

Figure 13-10. Steps for Changing Spanning-Tree Operation to the 802.1d Protocol

Enabling (or Disabling) Spanning Tree Operation on the Switch.

This command enables (or disables) spanning tree operation for either spanning tree version—STP/802.1d or RSTP/802.1w (the default). Before using this command, ensure that the version of spanning tree you want to use is active on the switch. (See the preceding topic, "Configuring the Switch To Use the 802.1d Spanning Tree Protocol (STP)" on page 13-23.)

Syntax: [no] spanning-tree

Default: Disabled

For example:

```
HPswitch spanning-tree
```

Enabling STP implements the spanning tree protocol for all physical ports on the switch, regardless of whether multiple VLANs are configured. Disabling STP removes protection against redundant loops that can significantly slow or halt a network.

This command enables STP with the current parameter settings or disables STP without losing the most-recently configured parameter settings. (To learn how the switch handles parameter changes, how to test changes without losing the previous settings, and how to replace previous settings with new settings, see Chapter 5, "Switch Memory and Configuration".) When enabling STP, you can also include the STP general and per-port parameters described in the next two sections. When you use the "no" form of the command, you can do so only to disable STP. (STP parameter settings are not changed when you disable STP.)

Caution

Because incorrect STP settings can adversely affect network performance, HP recommends that you use the default STP parameter settings. You should not change these settings unless you have a strong understanding of how STP operates. For more on STP, see the IEEE 802.1d standard.

Reconfiguring General STP Operation on the Switch. You can configure one or more of the following parameters:

Table 13-3.General STP Operating Parameters

Name	Default	Range	Function
priority	32768	0 - 65535	Specifies the priority value used along with the switch MAC address to determine which device is root. The lower a priority value, the higher the priority.
*maximum-age	20 seconds	6 - 40 seconds	Maximum received message age the switch allows for STP information before discarding the message.
*hello-time	2 seconds	1 - 10	Time between messages transmitted when the switch is the root.
*forward-delay	15 seconds	4 - 30 seconds	Time the switch waits before transitioning from the listening to the learning state, and between the learning state to the forwarding state.

*The switch uses its own maximum-age, hello-time, and forward-delay settings only if it is operating as the root device. If another device is operating as the root device, then the switch uses the other device's settings for these parameters.

Note

Executing **spanning-tree** alone enables STP. Executing spanning-tree with one or more of the above "STP Operating Parameters" does not enable STP. It only configures the STP parameters (regardless of whether STP is actually running (enabled) on the switch).

Syntax:

```
spanning-tree
    priority <0 - 65355>
    maximum-age <6 - 40 seconds>
    hello-time <1 - 10 seconds>
    forward-delay <4 - 30 seconds>
```

Default: See table 13-3, above.

For example, to configure a **maximum-age** of 30 seconds and a **hello-time** of 3 seconds for STP:

```
HPswitch(config)# spanning-tree maximum-age 30 hello-time
3
```

Reconfiguring Per-Port STP Operation on the Switch. This command enables STP (if not already enabled) and configures the following per-port parameters:

Table 13-4.Per-Port STP Parameters

Name	Default		Range	Function
path-cost	Ethernet: 10/100Tx: 100 Fx: Gigabit:	100 10 10 5	1 - 65535	Assigns an individual port cost that the switch uses to determine which ports are the forwarding ports.
priority	128		0 - 255	Used by STP to determine the port(s) to use for forwarding. The port with the lowest number has the highest priority.
mode	norm		norm *- or -* fast *- or -* uplink	Specifies whether a port progresses through the listening, learning, and forwarding (or blocking) states ("norm" mode) or transitions directly to the forwarding state ("fast" mode). • For information on when to use Fast mode, see "STP Fast Mode" on page 13-26.) • For information on Uplink mode, see "Fast-Uplink Spanning Tree Protocol (STP)" on page 13-27

You can also include STP general parameters in this command. See "Reconfiguring General STP Operation on the Switch" on page 13-24.

Syntax: spanning-tree [ethernet] <*port-list*>
path-cost <*1 - 65535*>
priority <*0 - 255*>
mode <*norm | fast*>

Default: See table 13-4, above.

For example, the following configures ports C5 and C6 to a path cost of **15**, a priority of **100**, and **fast** mode:

```
HPswitch(config)# spanning-tree c5-c6 path-cost 15
      priority 100 mode fast
```

STP Fast Mode

For standard STP operation, when a network connection is established on a device that is running STP, the port used for the connection goes through a sequence of states (Listening and Learning) before getting to its final state (Forwarding or Blocking, as determined by the STP negotiation). This sequence takes two times the forward delay value configured for the switch. The default is 15 seconds on HP switches, per the IEEE 802.1D standard recommendation, resulting in a total STP negotiation time of 30 seconds. Each switch port goes through this start-up sequence whenever the network connection is established on the port. This includes, for example, when the switch or connected device is powered up, or the network cable is connected.

A problem can arise from this long STP start-up sequence because some end nodes are configured to automatically try to access a network server whenever the end node detects a network connection. Typical server access includes to Novell servers, DHCP servers, and X terminal servers. If the server access is attempted during the time that the switch port is negotiating its STP state, the server access will fail. To provide support for this end node behavior, the Series 5300XL switches offers a configuration mode, called “Fast Mode”, that causes the switch port to skip the standard STP start-up sequence and put the port directly into the “Forwarding” state, thus allowing the server access request to be forwarded when the end node needs it.

If you encounter end nodes that repeatedly indicate server access failure when attempting to bring up their network connection, and you have enabled STP on the switch, try changing the configuration of the switch ports associated with those end nodes to STP Fast Mode.

Caution

The Fast Mode configuration should be used only on switch ports connected to end nodes. Changing the Mode to Fast on ports connected to hubs, switches, or routers may cause loops in your network that STP may not be able to immediately detect, in all cases. This will cause temporary loops in your network. After the fast start-up sequence, though, the switch ports operate according to the STP standard, and will adjust their state to eliminate continuing network loops.

To Enable or Disable Fast Mode for a Switch Port:

You can use either the CLI or the menu interface to toggle between STP Fast mode and STP Normal mode. (To use the menu interface, see “Menu: Configuring 802.1D STP” on page 13-19.)

Syntax: spanning-tree <*port list*> mode <fast | norm>

For example, to configure Fast mode for ports C1-C3 and C5:

```
HPswitch(config)# spanning-tree c1-c3,c5 mode fast
```

Fast-Uplink Spanning Tree Protocol (STP)

Fast-Uplink STP is an option added to the switch’s 802.1d STP to improve the recovery (convergence) time in wiring closet switches with redundant uplinks. Specifically, a Series 5300XL switch having redundant links toward the root device can decrease the convergence time (or failover) to a new uplink (STP root) port to as little as ten seconds. To realize this performance, the switch must be:

- Used as a wiring closet switch (also termed an *edge switch* or a *leaf switch*).
- Configured for fast-uplink STP mode on two or more ports intended for redundancy in the direction of the root switch, so that at any time only one of the redundant ports is expected to be in the forwarding state.

Note

Fast-Uplink STP operates only with 802.1d STP and is not available with the Rapid STP (802.1w) feature (page 13-7).

Caution

In general, fast-uplink spanning tree on the Series 5300XL switches is useful when running STP in a tiered topology that has well-defined edge switches. Also, ensure that an interior switch is used for the root switch and for any logical backup root switches. You can accomplish this by using the Spanning Tree Priority (sometimes termed bridge priority) settings that define the primary STP root switch and at least one failover root switch (in the event that the primary root switch fails). Inappropriate use of Fast-Uplink STP can cause intermittent loops in a network topology. For this reason, the Fast-Uplink STP feature should be used only by experienced network administrators who have a strong understanding of the IEEE 802.1D standard and STP interactions and operation. If you want to learn more about STP operation, you may find it helpful to refer to publications such as:

> Perlman, Radia, Interconnections, Second Edition; Bridges, Routers, Switches, and Internetworking Protocols, Addison-Wesley Professional Computing Series, October 1999

Note

When properly implemented, fast-uplink STP offers a method for achieving faster failover times than standard STP, and is intended for this purpose for instances where 802.1d STP has been chosen over 802.1w RSTP.

To use fast-uplink STP, configure fast-uplink (**Mode** = **Uplink**) only on the switch's upstream ports; (that is, two or more ports forming a group of redundant links in the direction of the STP root switch). If the active link in this group goes down, fast-uplink STP selects a different upstream port as the root port and resumes moving traffic in as little as ten seconds. The device(s) on the other end of the links must be running STP. However, because fast uplink should be configured only on the Series 5300XL switch uplink ports, the device(s) on the other end of the links can be either HP devices or another vendor's devices, regardless of whether they support fast uplink. For example:

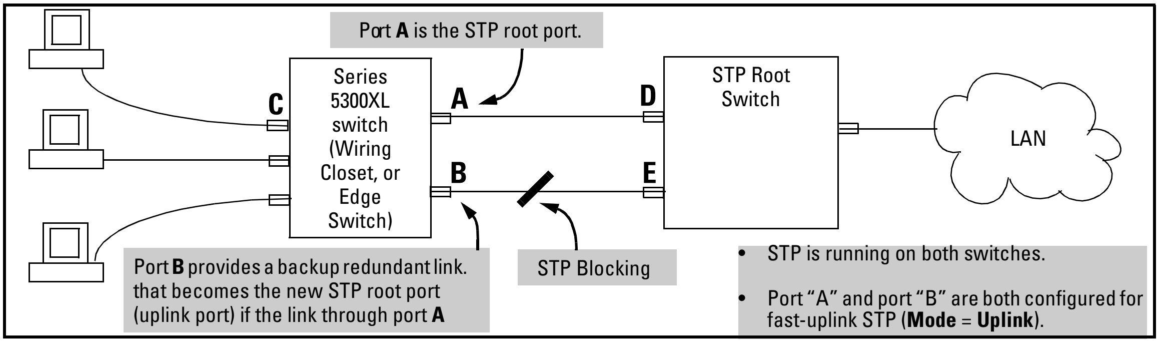
Port **A** is the STP root port.
Series 5300XL switch (Wiring Closet, or Edge Switch)
C
A
B
D
E
STP Root Switch
LAN
Port **B** provides a backup redundant link. that becomes the new STP root port (uplink port) if the link through port **A**
STP Blocking
• STP is running on both switches.
• Port "A" and port "B" are both configured for fast-uplink STP (**Mode** = **Uplink**).

Figure 13-11. Example of How To Implement Fast-Uplink STP

Terminology

Term	Definition
downlink port (downstream port)	A switch port that is linked to a port on another switch (or to an end node) that is sequentially further away from the STP root device. For example, port "C" in figure 13-11, above, is a downlink port.
edge switch	For the purposes of fast-uplink STP, this is a switch that has no other switches connected to its downlink ports. An edge switch is sequentially further from the root device than other switches to which it is connected. Also termed *wiring closet switch* or *leaf switch*. For example, switch "4" in figure 13-12 (page 29) is an edge switch.
interior switch	In an STP environment, a switch that is sequentially closer to the STP root device than one or more other switches to which it is connected. For example, switches "1", "2", and "3" in figure 13-12 (page 29) are interior switches.
single-instance spanning tree	A single spanning-tree ensuring that there are no logical network loops associated with any of the connections to the switch, regardless of whether there are any VLANs configured on the switch. For more information, see "Spanning Tree Protocol (STP)" in chapter 9, "Configuring Advanced Features", in the Management and Configuration Guide for your switch.
uplink port (upstream port)	A switch port linked to a port on another switch that is sequentially closer to the STP root device. For example, ports "A" and "B" in figure 13-11 on page 28 are uplink ports.
wiring closet switch	Another term for an "edge" or "leaf" switch.

When single-instance spanning tree (STP) is running in a network and a forwarding port goes down, a blocked port typically requires a period of

(2 x (*forward delay*) + link down detection)

to transition to forwarding. In a normal spanning tree environment, this transition is usually 30 seconds (with the **Forward Delay** parameter set to its default of 15 seconds). However, by using the fast-uplink spanning tree feature, a port on a Series 5300XL switch used as an *edge switch* can make this transition in as little as ten seconds. (In an STP environment, an *edge switch* is a switch that is connected only to switches that are closer to the STP root switch than the edge switch itself, as shown by switch "4" in figure 13-12, below.)

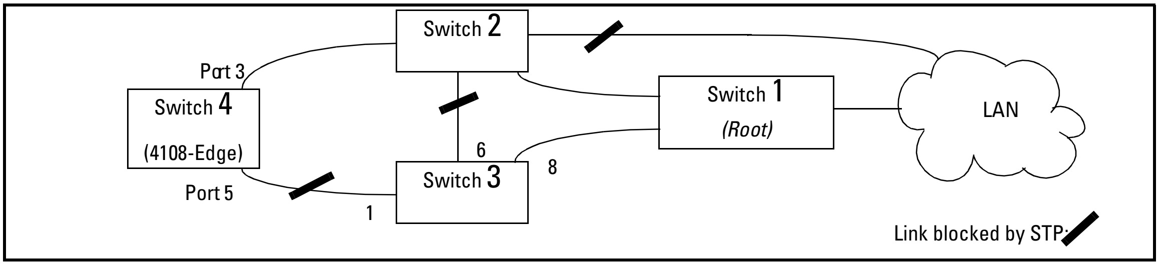

Figure 13-12. Example of an Edge Switch in a Topology Configured for STP Fast Uplink

In figure 13-12, STP is enabled and in its default configuration on all switches, unless otherwise indicated in table 13-13-5, below:

Table 13-5.STP Parameter Settings for Figure 13-12

STP Parameter	Switch “1”	Switch “2”	Switch “3”	Switch “4”
Switch Priority	0[1]	1[2]	32,768 (default)	32,768 (default)
(Fast) Uplink	No	No	No	Ports 3 & 5

[1] This setting ensures that Switch “1” will be the primary root switch for STP in figure 13-12.

[2] This setting ensures that Switch “2” will be the backup root switch for STP in figure 13-12.

With the above-indicated topology and configuration:

- **Scenario 1:** If the link between switches “4” and “2” goes down, then the link between switches “4” and “3” will begin forwarding in as little as ten seconds.
- **Scenario 2:** If Switch “1” fails, then:
  - Switch “2” becomes the root switch.
  - The link between Switch “3” and Switch “2” begins forwarding.
  - The link between Switch “2” and the LAN begins forwarding.

Operating Rules for Fast Uplink

- A switch with ports configured for fast uplink must be an edge switch and not either an interior switch or the STP root switch.

  Configure fast-uplink on only the edge switch ports used for providing redundant STP uplink connections in a network. (Configuring Fast-Uplink STP on ports in interior switches can create network performance problems.) That is, a port configured for STP uplink should not be connected to a switch that is sequentially further away from the STP root device. For example, switch “4” in figure 13-12 (page 13-29) is an edge switch.
- Configure fast uplink on a group (two or more) of redundant edge-switch uplink ports where only one port in the group is expected to be in the forwarding state at any given time.

- Edge switches cannot be directly linked together using fast-uplink ports. For example, the connection between switches 4 and 5 in figure 13-13 is not allowed for fast-uplink operation.

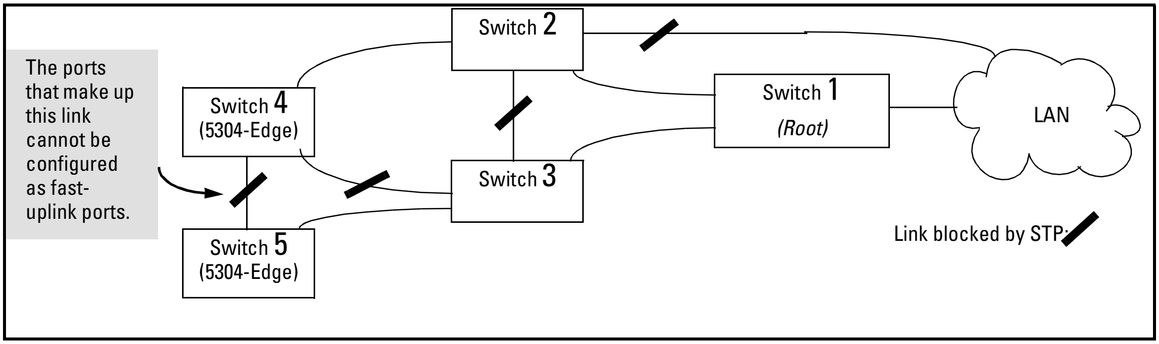

Figure 13-13. Example of a Disallowed Connection Between Edge Switches

- Apply fast-uplink only on the uplink ports of an edge switch. For example, on switch “4” (an edge switch) in figure 13-13 above, only the ports connecting switch “4” to switches “2” and “3” are upstream ports that would use fast uplink. Note also that fast uplink should *not* be configured on both ends of a point-to-point link, but only on the uplink port of an edge switch.
- Ensure that the switch you intend as a backup root device will in fact become the root if the primary root fails, and that no ports on the backup root device are configured for fast-uplink operation. For example, if the **STP Priority** is the same on all switches—default: 32768—then the switch with the lowest MAC address will become the root switch. If that switch fails, then the switch with the next-lowest MAC address will become the root switch. Thus, you can use **STP Priority** to control which switch STP selects as the root switch and which switch will become the root if the first switch fails.
- Fast-Uplink STP requires a minimum of two uplink ports.

Menu: Viewing and Configuring Fast-Uplink STP

You can use the menu to quickly display the entire STP configuration and to make any STP configuration changes.

To View and/or Configure Fast-Uplink STP. This procedure uses the Spanning Tree Operation screen to enable STP and to set the Mode for fast-uplink STP operation.

1. From the Main Menu select:

   2. Switch Configuration . . .

   4. Spanning Tree Operation

2. In the default STP configuration, RSTP is the selected protocol version. If this is the case on your switch, you must change the Protocol Version to STP in order to use Fast-Uplink STP:

- If the **Protocol Version** is set to RSTP (the default, as shown in this example, go to step 3.
- If the **Protocol Version** is set to STP, the rest of the screen will appear as shown in figure 13-16. In this case, go to step 4 on page 13-34.

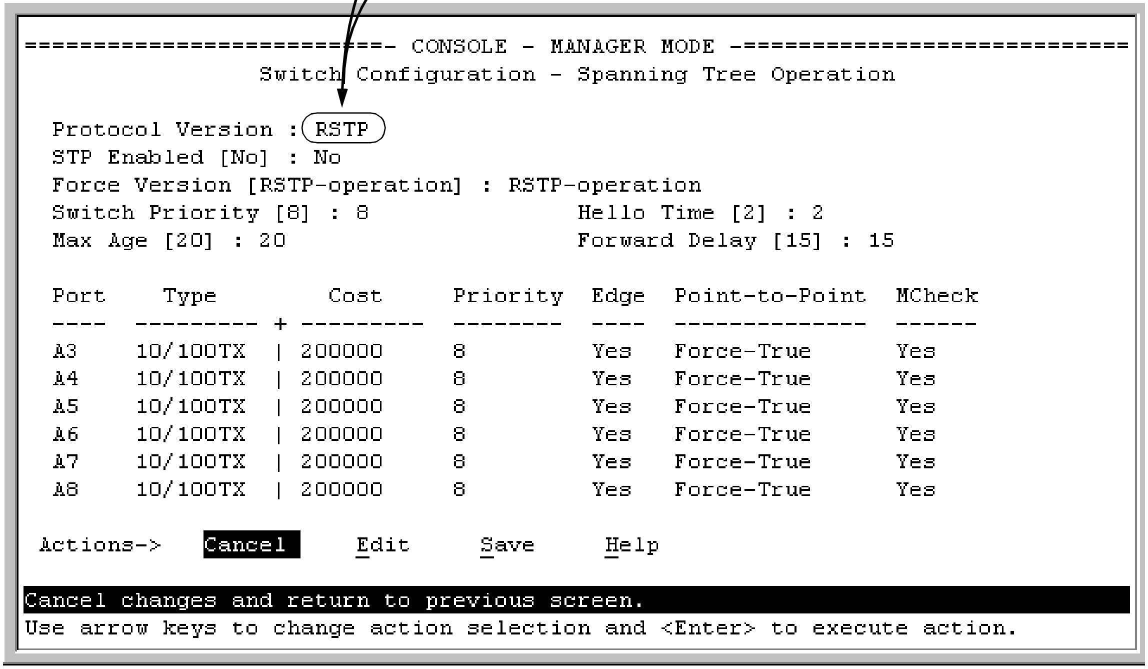

```
==========================- CONSOLE - MANAGER MODE -============================
                  Switch Configuration - Spanning Tree Operation

  Protocol Version : RSTP
  STP Enabled [No] : No
  Force Version [RSTP-operation] : RSTP-operation
  Switch Priority [8] : 8              Hello Time [2] : 2
  Max Age [20] : 20                    Forward Delay [15] : 15

  Port    Type        Cost      Priority  Edge  Point-to-Point  MCheck
  ----  --------- + ---------  --------  ----  --------------  ------
  A3    10/100TX  | 200000     8         Yes   Force-True      Yes
  A4    10/100TX  | 200000     8         Yes   Force-True      Yes
  A5    10/100TX  | 200000     8         Yes   Force-True      Yes
  A6    10/100TX  | 200000     8         Yes   Force-True      Yes
  A7    10/100TX  | 200000     8         Yes   Force-True      Yes
  A8    10/100TX  | 200000     8         Yes   Force-True      Yes

 Actions->  Cancel    Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 13-14. The Default STP Screen With the Protocol Version Field Set to "RSTP"

3. If the Protocol Version is set to RSTP (as shown in figure 13-14), do the following:

   a. Press [E] (**Edit**) to move the cursor to the **Protocol Version** field.

   b. Press the Space bar once to change the **Protocol Version** field to STP.

   c. Press [Enter] to return to the command line.

   d. Press [S] (for **Save**) to save the change and exit from the Spanning Tree Operation screen. you will then see a screen with the following:

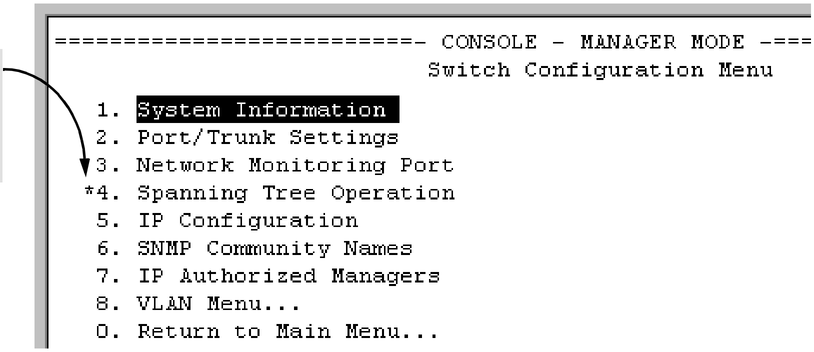

```
=========================- CONSOLE - MANAGER MODE -===
                                Switch Configuration Menu

  1. System Information
  2. Port/Trunk Settings
  3. Network Monitoring Port
 *4. Spanning Tree Operation
  5. IP Configuration
  6. SNMP Community Names
  7. IP Authorized Managers
  8. VLAN Menu...
  0. Return to Main Menu...
```

The asterisk indicates that you must reboot the switch to implement the configuration change from RSTP to STP.

Figure 13-15. Changing from RSTP to STP Requires a System Reboot

   e. Press [0] (zero) to return to the Main Menu, then [6] to reboot the switch.

   f. After you reboot the switch, enter the menu command at the CLI to return to the Main Menu, then select:

      2. Switch Configuration . . .
         4. Spanning Tree Operation

   You will then see the Spanning Tree screen with **STP** (802.1d) selected in the **Protocol Version** field (figure 13-16).

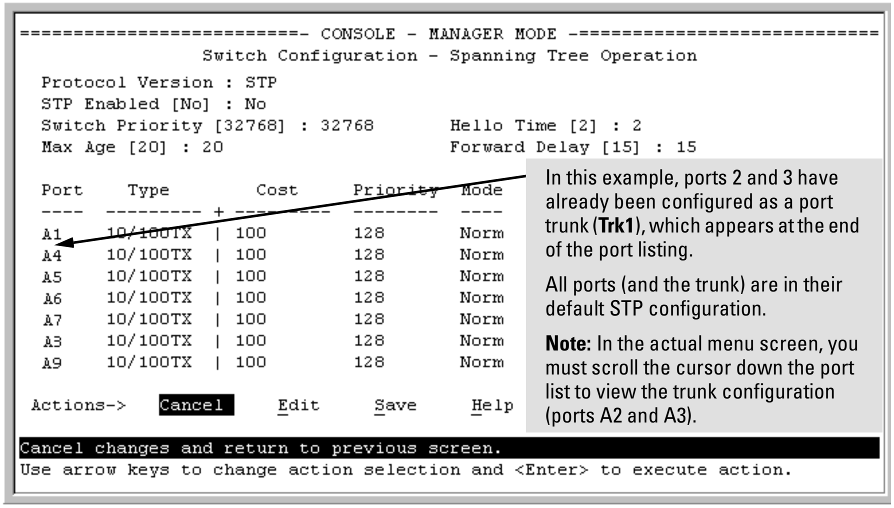

```
==========================- CONSOLE - MANAGER MODE -===========================
                 Switch Configuration - Spanning Tree Operation

  Protocol Version : STP
  STP Enabled [No] : No
  Switch Priority [32768] : 32768       Hello Time [2] : 2
  Max Age [20] : 20                     Forward Delay [15] : 15

  Port    Type        Cost      Priority  Mode
  ----  --------- + ---------  --------  ----
  A1    10/100TX  | 100        128       Norm
  A4    10/100TX  | 100        128       Norm
  A5    10/100TX  | 100        128       Norm
  A6    10/100TX  | 100        128       Norm
  A7    10/100TX  | 100        128       Norm
  A3    10/100TX  | 100        128       Norm
  A9    10/100TX  | 100        128       Norm

 Actions->  Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

In this example, ports 2 and 3 have already been configured as a port trunk (**Trk1**), which appears at the end of the port listing.

All ports (and the trunk) are in their default STP configuration.

Note: In the actual menu screen, you must scroll the cursor down the port list to view the trunk configuration (ports A2 and A3).

Figure 13-16. The Spanning Tree Operation Screen

4. On the ports and/or trunks you want to use for redundant fast uplink connections, change the mode to **Uplink**. In this example, port A1 and Trk1 (using ports A2 and A3) provide the redundant uplinks for STP:

   a. Press [E] (for **Edit**), then enable STP on the switch by using the Space bar to select **Yes** in the Spanning Tree Enabled field.

   b. Use [Tab] to move to the Mode field for port A1.

   c. Use the Space bar to select **Uplink** as the mode for port A1.

   d. Use [↓] to move to the Mode field for Trk1.

   e. Use the Space bar to select **Uplink** as the Mode for Trk1.

   f. Press [Enter] to return the cursor to the Actions line.

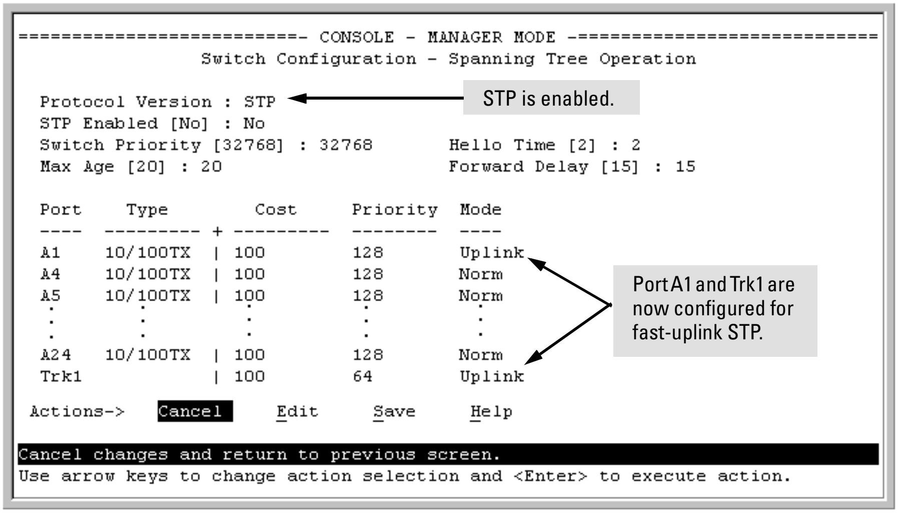

```
=========================- CONSOLE - MANAGER MODE -===========================
                 Switch Configuration - Spanning Tree Operation

  Protocol Version : STP
  STP Enabled [No] : No
  Switch Priority [32768] : 32768       Hello Time [2] : 2
  Max Age [20] : 20                     Forward Delay [15] : 15

  Port    Type        Cost      Priority  Mode
  ----   --------- + ---------  --------  ----
  A1     10/100TX  | 100        128       Uplink
  A4     10/100TX  | 100        128       Norm
  A5     10/100TX  | 100        128       Norm
  .       .          .          .          .
  .       .          .          .          .
  A24    10/100TX  | 100        128       Norm
  Trk1             | 100        64        Uplink

 Actions->   Cancel     Edit      Save      Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

STP is enabled.

Port A1 and Trk1 are now configured for fast-uplink STP.

Figure 13-17. Example of STP Enabled with Two Redundant Links Configured for Fast-Uplink STP

5. Press [S] (for **Save**) to save the configuration changes to flash (non-volatile) memory.

To View Fast-Uplink STP Status. Continuing from figures 13-16 and 13-17 in the preceding procedure, this task uses the same screen that you would use to view STP status for other operating modes.

1. From the Main Menu, select:

1. Status and Counters . . .

7. Spanning Tree Information

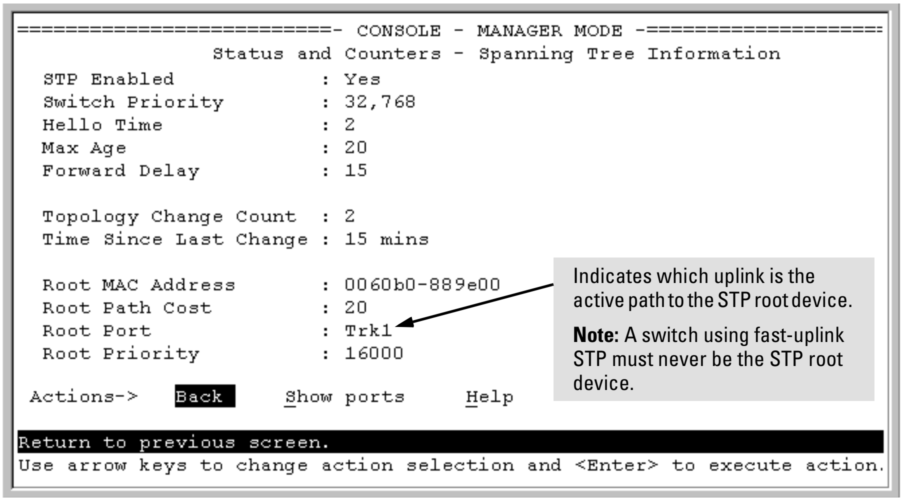

```
==========================- CONSOLE - MANAGER MODE -====================
                 Status and Counters - Spanning Tree Information
  STP Enabled              : Yes
  Switch Priority          : 32,768
  Hello Time               : 2
  Max Age                  : 20
  Forward Delay            : 15

  Topology Change Count    : 2
  Time Since Last Change   : 15 mins

  Root MAC Address         : 0060b0-889e00
  Root Path Cost           : 20
  Root Port                : Trk1
  Root Priority            : 16000

 Actions->   Back     Show ports     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Indicates which uplink is the active path to the STP root device.

Note: A switch using fast-uplink STP must never be the STP root device.

Figure 13-18. Example of STP Status with Trk1 (Trunk 1) as the Path to the STP Root Device

2. Press S (for **Show ports**) to display the status of individual ports.

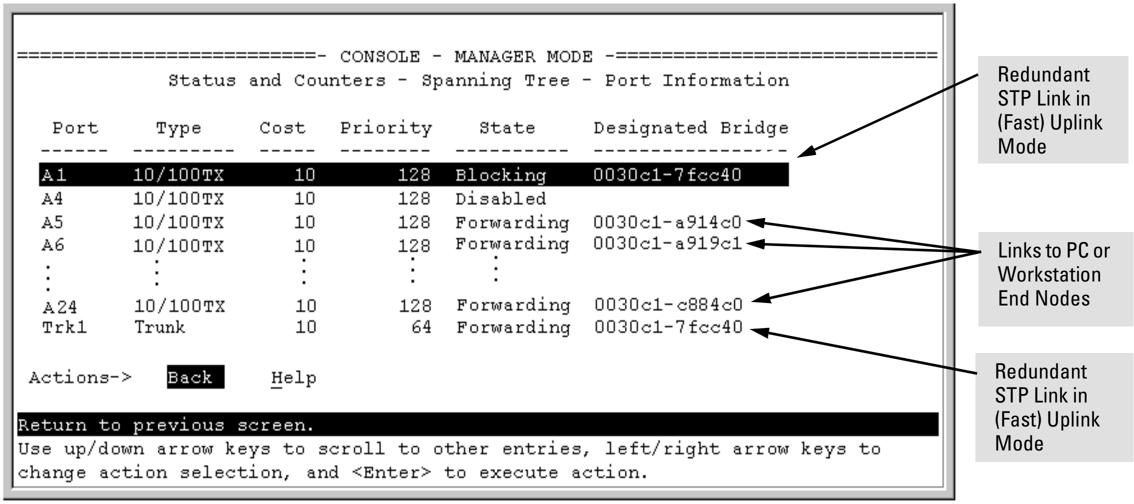

```
==========================- CONSOLE - MANAGER MODE -=============================
                 Status and Counters - Spanning Tree - Port Information

  Port     Type       Cost   Priority    State       Designated Bridge
 ------  ---------   -----  --------   ----------   -----------------
 A1      10/100TX      10       128   Blocking     0030c1-7fcc40
 A4      10/100TX      10       128   Disabled
 A5      10/100TX      10       128   Forwarding   0030c1-a914c0
 A6      10/100TX      10       128   Forwarding   0030c1-a919c1
 :          :           :         :       :
 A24     10/100TX      10       128   Forwarding   0030c1-c884c0
 Trk1    Trunk         10        64   Forwarding   0030c1-7fcc40

 Actions->   Back     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Redundant STP Link in (Fast) Uplink Mode

Links to PC or Workstation End Nodes

Redundant STP Link in (Fast) Uplink Mode

Figure 13-19. Example of STP Port Status with Two Redundant STP Links

In figure 13-19:

- Port A1 and Trk1 (trunk 1; formed from ports 2 and 3) are redundant fast-uplink STP links, with trunk 1 forwarding (the active link) and port A1 blocking (the backup link). (To view the configuration for port A1 and Trk1, see figure 13-17 on page 13-35.)
- If the link provided by trunk 1 fails (on both ports), then port A1 begins forwarding in fast-uplink STP mode.
- Ports A5, A6, and A24 are connected to end nodes and do not form redundant links.

CLI: Viewing and Configuring Fast-Uplink STP

Using the CLI to View Fast-Uplink STP. You can view fast-uplink STP using the same **show** commands that you would use for standard STP operation:

Syntax: `show spanning-tree` Lists STP status.
`show spanning-tree config` Lists STP configuration for the switch and for individual ports.

For example, figures 13-20 and 13-21 illustrate a possible topology, STP status listing, and STP configuration for a Series 5300XL switch with:

- STP enabled and the switch operating as an Edge switch
- Port A1 and trunk 1 (Trk1) configured for fast-uplink STP operation
- Several other ports connected to PC or workstation end nodes

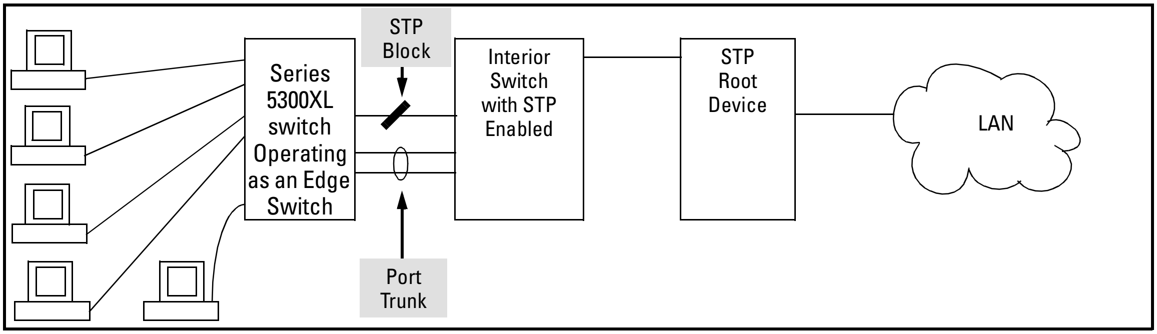
STP Block
Series 5300XL switch Operating as an Edge Switch
Interior Switch with STP Enabled
STP Root Device
LAN
Port Trunk

Figure 13-20. Example Topology for the Listing Shown in Figure 13-21

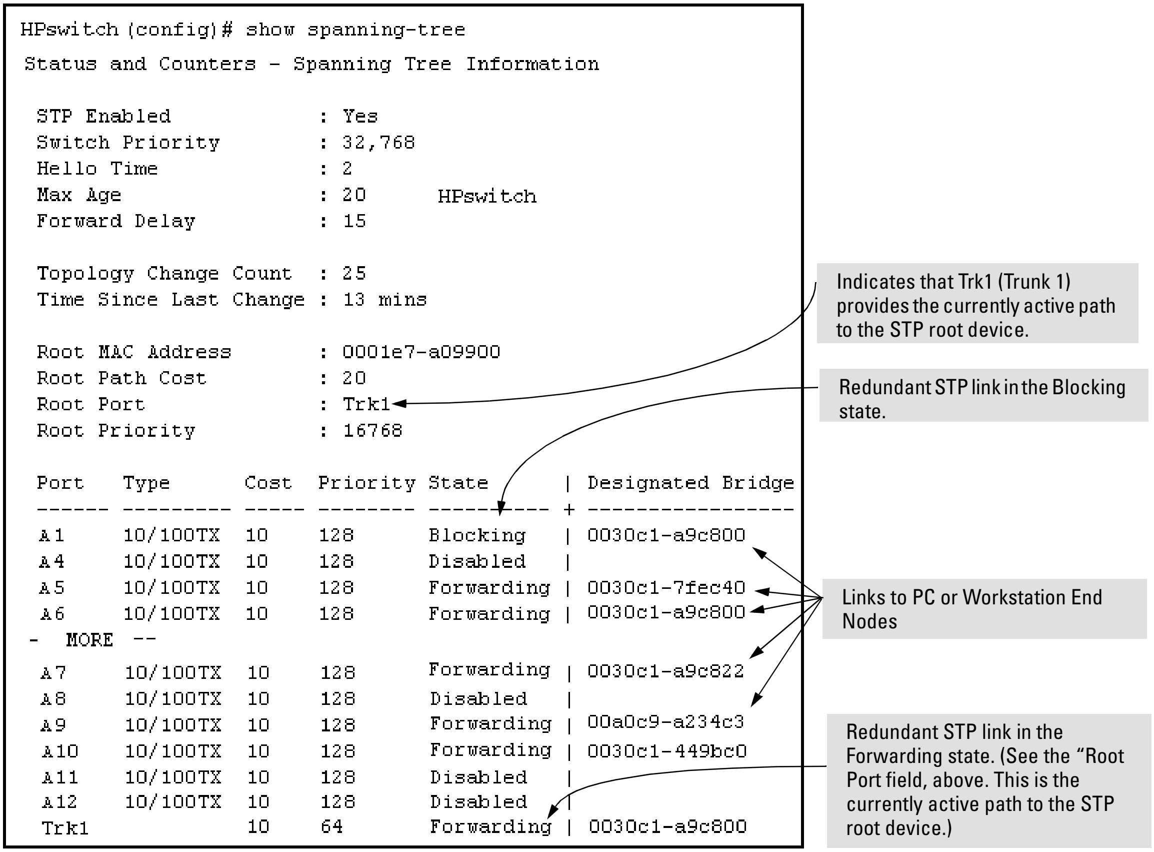

```
HPswitch(config)# show spanning-tree
 Status and Counters - Spanning Tree Information

  STP Enabled            : Yes
  Switch Priority        : 32,768
  Hello Time             : 2
  Max Age                : 20      HPswitch
  Forward Delay          : 15

  Topology Change Count  : 25
  Time Since Last Change : 13 mins

  Root MAC Address       : 0001e7-a09900
  Root Path Cost         : 20
  Root Port              : Trk1
  Root Priority          : 16768

  Port   Type      Cost  Priority State      | Designated Bridge
  ------ --------- ----- -------- ---------- + -----------------
  A1     10/100TX  10    128      Blocking   | 0030c1-a9c800
  A4     10/100TX  10    128      Disabled   |
  A5     10/100TX  10    128      Forwarding | 0030c1-7fec40
  A6     10/100TX  10    128      Forwarding | 0030c1-a9c800
 -  MORE  --
  A7     10/100TX  10    128      Forwarding | 0030c1-a9c822
  A8     10/100TX  10    128      Disabled   |
  A9     10/100TX  10    128      Forwarding | 00a0c9-a234c3
  A10    10/100TX  10    128      Forwarding | 0030c1-449bc0
  A11    10/100TX  10    128      Disabled   |
  A12    10/100TX  10    128      Disabled   |
  Trk1             10    64       Forwarding | 0030c1-a9c800
```

Indicates that Trk1 (Trunk 1) provides the currently active path to the STP root device.

Redundant STP link in the Blocking state.

Links to PC or Workstation End Nodes

Redundant STP link in the Forwarding state. (See the "Root Port field, above. This is the currently active path to the STP root device.)

Figure 13-21. Example of a Show Spanning-Tree Listing for the Topology Shown in Figure 13-20

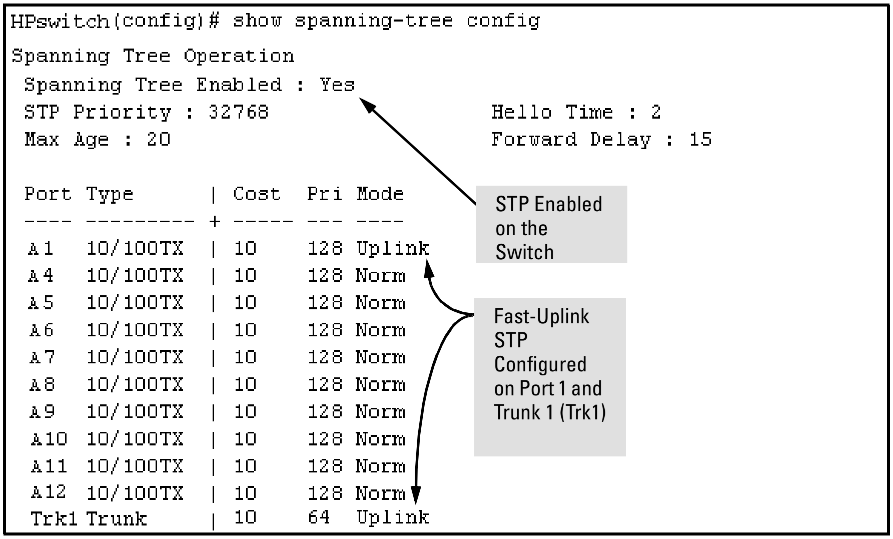

```
HPswitch(config)# show spanning-tree config

Spanning Tree Operation
 Spanning Tree Enabled : Yes
 STP Priority : 32768                    Hello Time : 2
 Max Age : 20                            Forward Delay : 15

 Port Type      | Cost  Pri Mode
 ---- --------- + ----- --- ----
  A1  10/100TX  | 10    128 Uplink
  A4  10/100TX  | 10    128 Norm
  A5  10/100TX  | 10    128 Norm
  A6  10/100TX  | 10    128 Norm
  A7  10/100TX  | 10    128 Norm
  A8  10/100TX  | 10    128 Norm
  A9  10/100TX  | 10    128 Norm
  A10 10/100TX  | 10    128 Norm
  A11 10/100TX  | 10    128 Norm
  A12 10/100TX  | 10    128 Norm
  Trk1 Trunk    | 10    64  Uplink
```

STP Enabled on the Switch

Fast-Uplink STP Configured on Port 1 and Trunk 1 (Trk1)

Figure 13-22. Example of a Configuration Supporting the STP Topology Shown in Figure 13-20

Using the CLI To Configure Fast-Uplink STP. This example uses the CLI to configure the switch for the fast-uplink operation shown in figures 13-20, 13-21, and 13-22. (The example assumes that ports A2 and A3 are already configured as members of the port trunk—Trk1, and all other STP parameters are left in their default state.)

Note that the default STP Protocol Version is RSTP (Rapid STP, or 802.1w). Thus, if the switch is set to the STP default, you must change it to the STP (802.1d) Protocol Version before you can configure Fast-Uplink. For example:

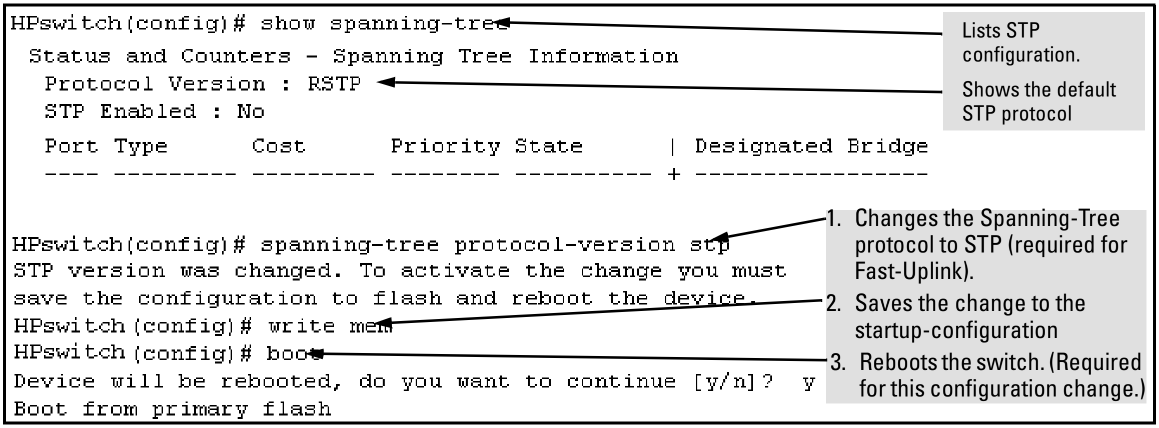

```
HPswitch(config)# show spanning-tree
 Status and Counters - Spanning Tree Information
  Protocol Version : RSTP
  STP Enabled : No
  Port Type      Cost      Priority State      | Designated Bridge
  ---- --------- --------- -------- ---------- + ----------------


HPswitch(config)# spanning-tree protocol-version stp
STP version was changed. To activate the change you must
save the configuration to flash and reboot the device.
HPswitch(config)# write mem
HPswitch(config)# boot
Device will be rebooted, do you want to continue [y/n]?  y
Boot from primary flash
```

Lists STP configuration.

Shows the default STP protocol

1. Changes the Spanning-Tree protocol to STP (required for Fast-Uplink).
2. Saves the change to the startup-configuration
3. Reboots the switch. (Required for this configuration change.)

Figure 13-23. Example of Changing the STP Configuration from the Default RSTP (802.1w) to STP (802.1d)

Syntax:spanning-tree e <*port/trunk-list*> mode uplinkEnables STP on the switch and configures
fast-uplink STP on the designated
interfaces (port or trunk).

For example:

```
HPswitch(config)# spanning-tree e A1,trk1 mode uplink
```

Operating Notes

Effect of Reboots on Fast-Uplink STP Operation. When configured, fast-uplink STP operates on the designated ports in a running switch. However, if the switch experiences a reboot, the fast-uplink ports (Mode = **Uplink**) use the longer forwarding delay used by ports on standard 802.1D STP (non fast-uplink). This prevents temporary loops that could otherwise result while the switch is determining the STP status for all ports. That is, on ports configured for fast-uplink STP, the first STP state transition after a reboot takes the same amount of time as for redundant ports that are not configured for fast-uplink STP.

Using Fast Uplink with Port Trunks. To use a port trunk for fast-uplink STP, configure it in the same way that you would an individual port for the same purpose. A port trunk configured for fast uplink operates in the same way as an individual, non-trunked port operates; that is, as a logical port.

Note

When you add a port to a trunk, the port takes on the STP mode configured for the trunk, regardless of which STP mode was configured on the port before it was added to the trunk. Thus, all ports belonging to a trunk configured with **Uplink** in the STP **Mode** field will operate in the fast-uplink mode. (If you remove a port from a trunk, the port reverts to the STP Mode setting it had before you added the port to the trunk.

To use fast uplink over a trunk, you must:

1. Create the trunk.
2. Configure the trunk for fast uplink in the same way that you would configure an individual port for fast uplink.

When you first create a port trunk, its STP Mode setting will be **Norm**, regardless of whether one or more ports in the trunk are set to fast uplink (Mode = **Uplink**). You must still specifically configure the trunk Mode setting to **Uplink**. Similarly, if you eliminate a trunk, the Mode setting on the individual ports in the trunk will return to their previous settings.

For Troubleshooting Information on Fast Uplink. Refer to “Spanning-Tree Protocol (STP) and Fast-Uplink Problems” on page C-15 (in the “Troubleshooting” appendix).

Web: Enabling or Disabling STP

In the web browser interface you can enable or disable STP on the switch. To configure other STP features, telnet to the switch console and use the CLI.

To enable or disable STP on the switch:

1. Click on the **Configuration** tab
2. Click on [Device Features].
3. Enable or disable STP.
4. Click on [Apply Changes] to implement the configuration change.

For web-based help on how to use the web browser interface screen, click on the [?] button provided on the web browser screen.

14

Switch Meshing

Contents

Introduction

Switch meshing is a load-balancing technology that enhances reliability and performance in these ways:

- Provides significantly better bandwidth utilization than either Spanning Tree Protocol (STP) or standard port trunking.
- Uses redundant links that remain open to carry traffic, removing any single point of failure for disabling the network, and allowing quick responses to individual link failures. This also helps to maximize investments in ports and cabling.
- Unlike trunked ports, the ports in a switch mesh can be of different types and speeds. For example, a 10Base-FL port and a 1GB port can be included in the same switch mesh.

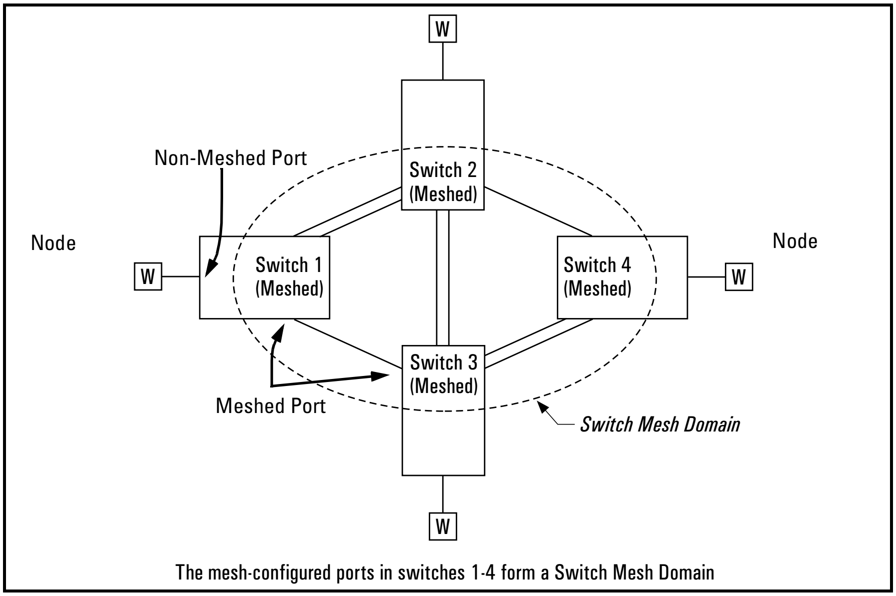

Figure 14-1. Example of Switch Meshing

Finding the Fastest Path. Using multiple switches redundantly linked together to form a *meshed switch domain*, switch meshing dynamically distributes traffic across load-balanced switch paths by seeking the fastest paths for new traffic between nodes. In actual operation, the switch mesh periodically determines the best (lowest latency) paths, then assigns these paths as the need arises. The path assignment remains until the related MAC address entry times out. The mesh sees later traffic between the same nodes as new traffic, and may assign a different path, depending on conditions at the time. For example, at one time the best path from node A to node B is through switch 2. However, if traffic between node A and node B ceases long enough for the path assignment to age out, then the next time node A has traffic for node B, the assigned path between these nodes may be through switch 3 if network conditions have changed significantly.

Note

The **mac-age-time** parameter determines how long an inactive path assignment remains in memory. Refer to "Reconfigure the Age Time for Learned MAC Addresses" on page 6-13.

Because Redundant Paths Are Active, Meshing Adjusts Quickly to Link Failures. If a link in the mesh fails, the fast convergence time designed into meshing typically has an alternate route selected in less than a second for traffic that was destined for the failed link.

Meshing Allows Scalable Responses to Increasing Bandwidth Demand. As more bandwidth is needed in a LAN backbone, another switch and another set of links can be added. This means that bandwidth is not limited by the number of trunk ports allowed in a single switch.

Switch Meshing Fundamentals

Terminology

Switch Mesh Domain. This is a group of meshed switch ports exchanging meshing protocol packets. Paths between these ports can have multiple redundant links without creating broadcast storms.

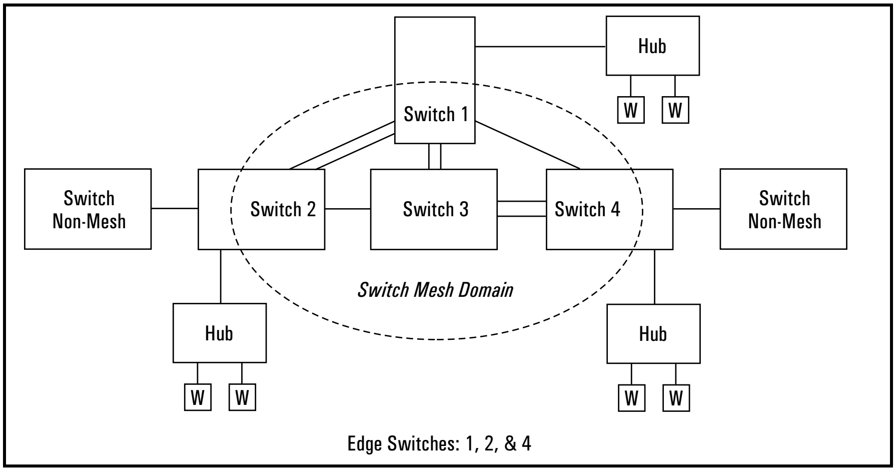

Figure 14-2. Example of a Switch Mesh Domain in a Network

Edge Switch. This is a switch that has some ports in the switch meshing domain and some ports outside of the domain. (See figure 14-2, above.)

Operating Rules

(See also “Requirements and Restrictions” on page 14-24.)

- A meshed switch can have some ports in the meshed domain and others outside the meshed domain. That is, ports within the meshed domain must be configured for meshing, while ports outside the meshed domain must not be configured for meshing.
- Meshed links must be point-to-point switch links.
- Within any meshed switch, all ports belong to the same meshed domain.
- A switch can have up to 24 meshed ports.
- A mesh domain can include up to 12 switches.
- Hub links between meshed switch links are not allowed.

Note

Switch mesh domains do not allow inclusion of:

- Hubs
- Switches that are not configured for meshing
- Non-meshed ports in a switch that has some meshed ports

Linking a non-mesh device or port into the mesh causes the meshed switch port(s) connected to that device to shut down.

- If the switch has multiple static VLANs and you configure a port for meshing, the port becomes a tagged member of all such VLANs. If you remove a port from meshing, it becomes an untagged member of only the default VLAN.
- A port configured as a member of a static trunk (LACP, FEC, or Trunk type) cannot be configured for meshing, and the reverse.
- A port with 802.1x port access security is not supported for switch meshing.
- If a port belongs to a dynamic LACP trunk, you can impose meshing on the port without having to remove the port from the trunk. (If a port is a member of a dynamic LACP trunk and you configure the port for meshing, it automatically ceases to be a member of the dynamic trunk.)
- On a port configured for meshing, if you subsequently remove meshing from the port’s configuration and reboot the switch, the port returns to its default configuration. (It *does not* revert to any non-default configuration it had before being configured for meshing).
- When meshing is configured on the switch, the routing features (IP routing, RIP, and OSPF) must be disabled, and the reverse. *That is, the switch’s meshing and routing features cannot be enabled at the same time.*

- Spanning tree must be the same for all switches in the mesh (enabled or disabled). If spanning tree is enabled in the mesh, it must be the same version on all switches in the mesh (STP or RSTP).
- If a switch in the mesh has GVRP enabled, then all switches in the mesh must have GVRP enabled. Otherwise, traffic on a dynamic VLAN may not pass through the mesh.

GVRP Note

HP Procurve 1600M/2400M/2424M/4000M/8000M switches do not offer the GVRP feature. If any of these switches are in your switch mesh, then GVRP must be disabled on any Series 5300XL switches in the mesh.

- If GVRP is enabled, meshed ports in a switch become members of any dynamic VLANs created in the switch in the same way that they would if meshing was not configured in the switch. (For more on GVRP, refer to "GVRP" on page 11-34.)
- If a switch in the mesh has a particular static vlan enabled, then all switches in the mesh must have that static vlan enabled.
- If a switch in the mesh has IGMP enabled, then all switches in the mesh must have IGMP enabled.
- After adding or removing a port from the mesh, you must save the current configuration and reboot the switch in order for the change to take effect.
- Multiple meshed domains require separation by either a non-meshed switch or a non-meshed link. For example:

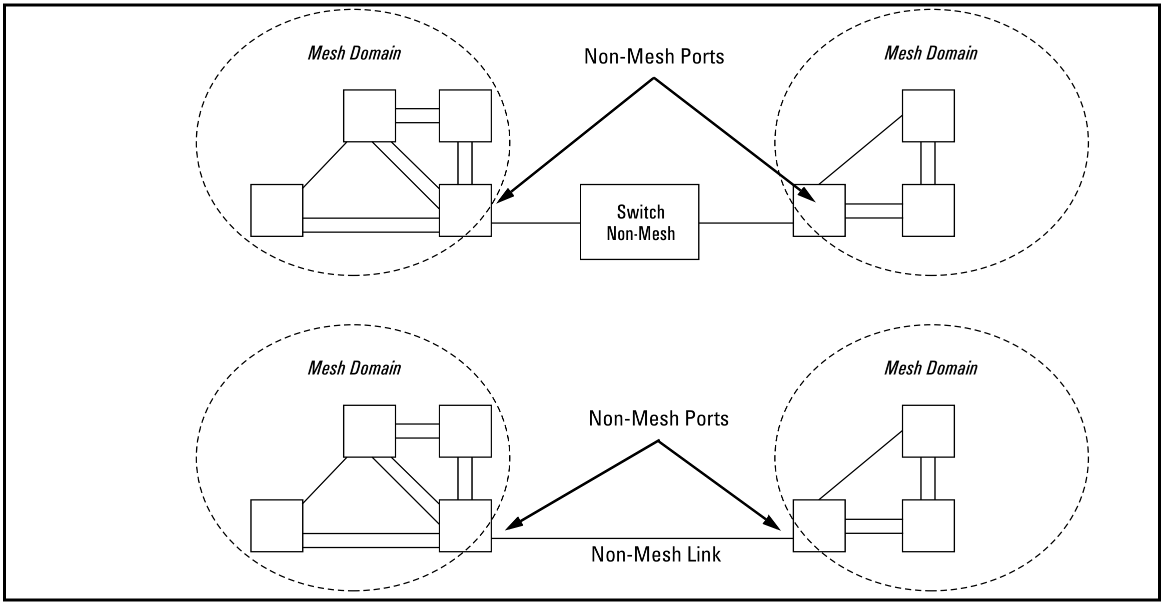

Figure 14-3. Example of Multiple Meshed Domains Separated by a Non-Mesh Switch or a Non-Mesh Link

Using a Heterogeneous Switch Mesh

You can use Series 5300XL switches together with any of the older HP Procurve Switch 1600M/2400M/2424M/4000M/8000M models. Note that if you connect one of these older models to a Series 5300XL switch in a mesh environment, up to two minutes may elapse before both devices become established in the mesh. These restrictions also apply:

- The older models cannot be used in a mesh environment with Series 5300XL switches where there is a duplicate MAC address on multiple switches and different VLANs. If you add an older model switch in this environment after the mesh is established, this switch will not be admitted to the mesh. If an older model switch is operating in a mesh with Series 5300XL switches and you introduce a topology that creates a duplicate MAC address on multiple switches, the device accessed by these multiple switches will be blocked. For example:

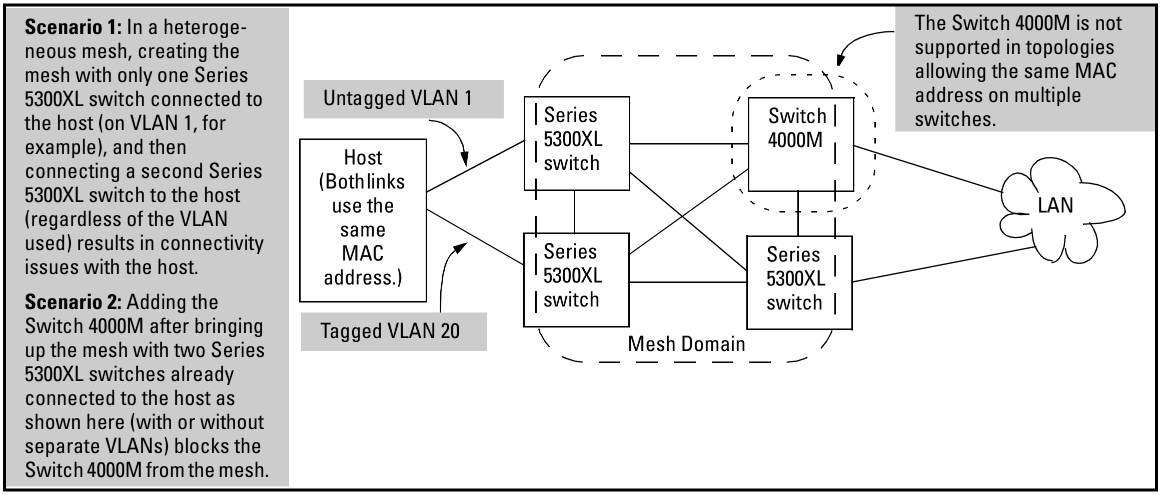

Scenario 1: In a heterogeneous mesh, creating the mesh with only one Series 5300XL switch connected to the host (on VLAN 1, for example), and then connecting a second Series 5300XL switch to the host (regardless of the VLAN used) results in connectivity issues with the host.

Scenario 2: Adding the Switch 4000M after bringing up the mesh with two Series 5300XL switches already connected to the host as shown here (with or without separate VLANs) blocks the Switch 4000M from the mesh.

The Switch 4000M is not supported in topologies allowing the same MAC address on multiple switches.

Figure 14-4. Example of an Unsupported Heterogeneous Topology Where Duplicate MAC Addresses Come Through Different Switches (Regardless of the VLANs Used)

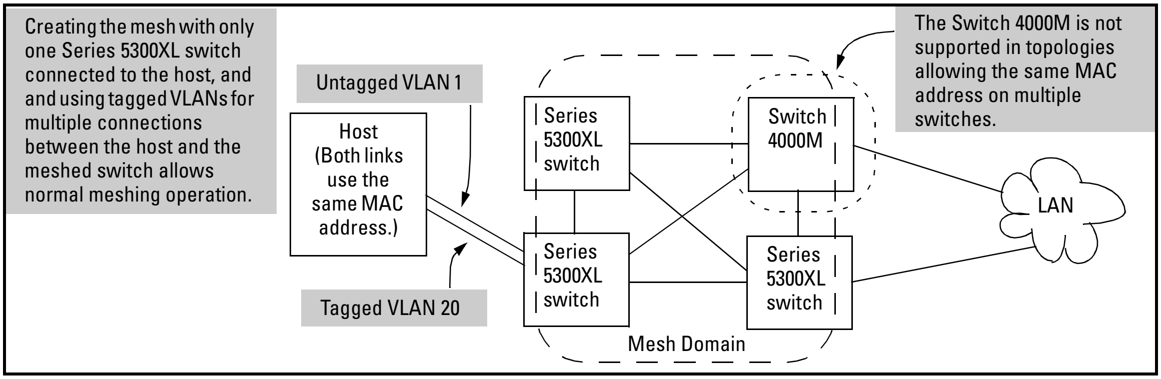

Creating the mesh with only one Series 5300XL switch connected to the host, and and using tagged VLANs for multiple connections between the host and the meshed switch allows normal meshing operation.

The Switch 4000M is not supported in topologies allowing the same MAC address on multiple switches.

Figure 14-5. Example of a Supported Heterogeneous Topology Where Duplicate MAC Addresses Come Through Different VLANs on the Same Switch

Note that in figures 14-4 and 14-5, if all switches are Series 5300XL switches, then you can use either topology.

Also, if you have two separate switch meshes with the topology shown in figure 14-6, you cannot join them into a single mesh.

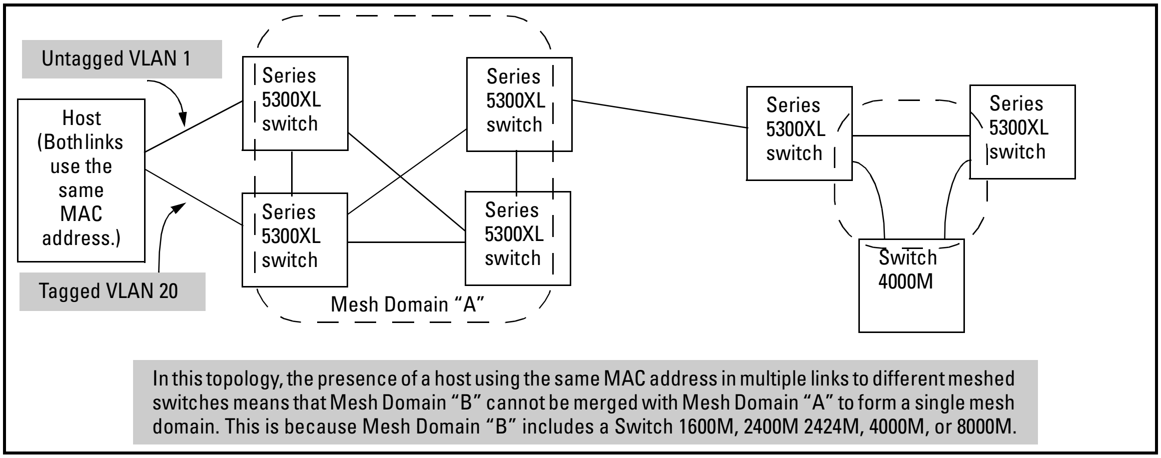

In this topology, the presence of a host using the same MAC address in multiple links to different meshed switches means that Mesh Domain “B” cannot be merged with Mesh Domain “A” to form a single mesh domain. This is because Mesh Domain “B” includes a Switch 1600M, 2400M 2424M, 4000M, or 8000M.

Figure 14-6. Example of Topology Where Adjacent Switch Meshes Cannot Be Merged Into a Single Mesh

- If you remove an older model HP Procurve switch from a mesh and leave only Series 5300XL switches in the mesh, it can take up to two minutes before the Series 5300XL switches begin accepting duplicate MAC addresses on different Switches.
- Automatic Broadcast Control (ABC) on HP Procurve 8000M/4000M/2424M/2400M/1600M switches is not supported when these switches are used in the same mesh domain with Series 5300XL switches. Thus, in a mesh domain populated with both types of switches, ABC must be disabled (the default setting) on all of the 8000M/4000M/2424M/2400M/1600M switches in the domain.

Bringing Up a Switch Mesh Domain:

When a meshed port detects a non-meshed port on the opposite end of a point-to-point connection, the link will be blocked. Thus, as you bring up switch meshing on various switches, you may temporarily experience blocked ports where meshed links should be running. These conditions should clear themselves after all switches in the mesh have been configured for meshing and their switches rebooted. To reduce the effect of blocked ports during bring-up, configure meshing and reboot the switches before installing the meshed switches in the network.

Further Operating Information

Refer to “Operating Notes for Switch Meshing” on page 14-17.

Configuring Switch Meshing

Preparation

Before configuring switch meshing:

- Review the Operating Rules (page 14-5), and particularly the restrictions and requirements for using switch meshing in environments that include static trunks, multiple static VLANs, GVRP, IGMP, and STP.
- To avoid unnecessary system disruption, plan the mesh bring-up to minimize temporary port-blocking. (Refer to “Bringing Up a Switch Mesh Domain:” on page 14-9.)
- Note
- To view the current switch mesh status on the Series 5300XL switches, use the CLI **show mesh** command (page 14-14).

Menu: To Configure Switch Meshing

1. From the Main Menu, select:

   2. Switch Configuration
   2. Port/Trunk Settings

2. Press [E] (for **Edit**) to access the load balancing parameters.

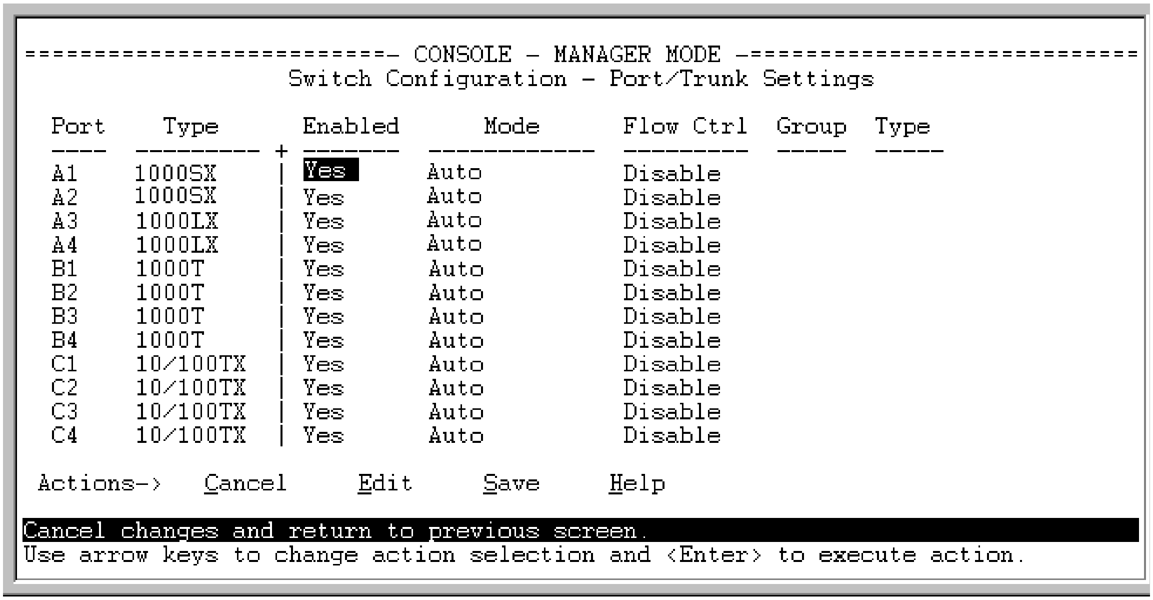

```
========================- CONSOLE - MANAGER MODE -============================
                     Switch Configuration - Port/Trunk Settings

  Port    Type       Enabled      Mode       Flow Ctrl  Group  Type
  ----  ---------- + -------  -----------  ---------  -----  -----
  A1    1000SX     | Yes      Auto         Disable
  A2    1000SX     | Yes      Auto         Disable
  A3    1000LX     | Yes      Auto         Disable
  A4    1000LX     | Yes      Auto         Disable
  B1    1000T      | Yes      Auto         Disable
  B2    1000T      | Yes      Auto         Disable
  B3    1000T      | Yes      Auto         Disable
  B4    1000T      | Yes      Auto         Disable
  C1    10/100TX   | Yes      Auto         Disable
  C2    10/100TX   | Yes      Auto         Disable
  C3    10/100TX   | Yes      Auto         Disable
  C4    10/100TX   | Yes      Auto         Disable

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 14-7. Example of the Screen for Configuring Ports for Meshing

3. In the Group column, move the cursor to the port you want to assign to the switch mesh.

4. Press [M] to choose **Mesh** for the selected port.

5. Use the [↑] or the [↓] key to select the next port you want to include in your mesh domain, then press [M] again. For example, if you were adding ports A1 and A2 to your mesh domain, the screen would appear similar to figure 14-8:

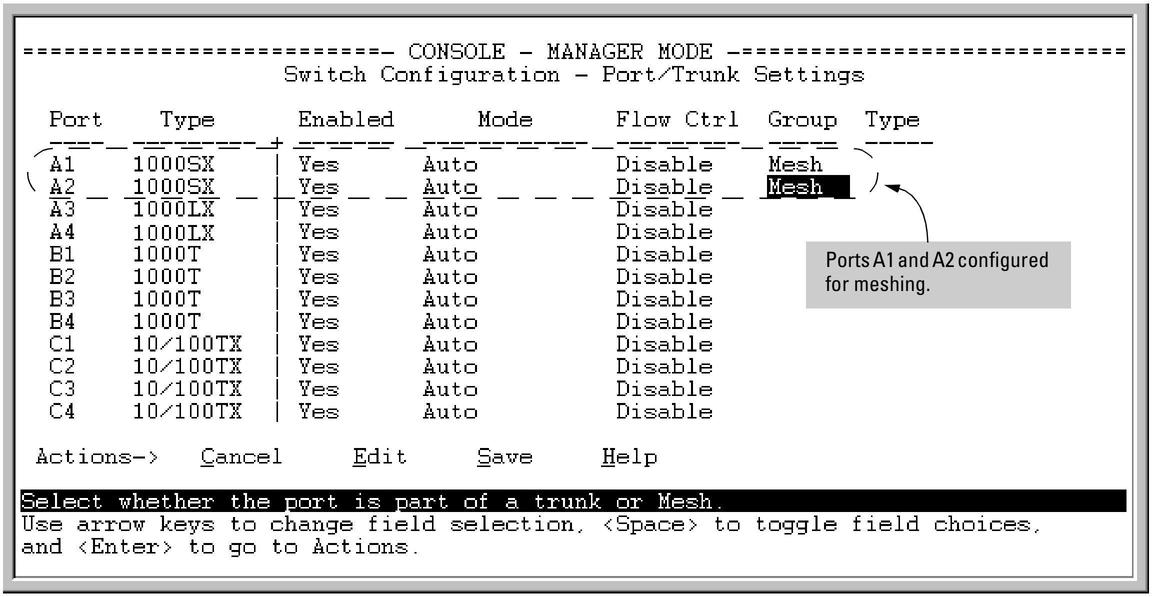

```
==========================- CONSOLE - MANAGER MODE -============================
                  Switch Configuration - Port/Trunk Settings

  Port    Type      Enabled       Mode       Flow Ctrl  Group  Type
  ----  ---------+ -------  ------------  ---------  -----  -----
  A1    1000SX    | Yes      Auto          Disable    Mesh
  A2    1000SX    | Yes      Auto          Disable    Mesh
  A3    1000LX    | Yes      Auto          Disable
  A4    1000LX    | Yes      Auto          Disable
  B1    1000T     | Yes      Auto          Disable
  B2    1000T     | Yes      Auto          Disable
  B3    1000T     | Yes      Auto          Disable
  B4    1000T     | Yes      Auto          Disable
  C1    10/100TX  | Yes      Auto          Disable
  C2    10/100TX  | Yes      Auto          Disable
  C3    10/100TX  | Yes      Auto          Disable
  C4    10/100TX  | Yes      Auto          Disable

 Actions->   Cancel     Edit     Save     Help

Select whether the port is part of a trunk or Mesh.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Ports A1 and A2 configured for meshing.

Figure 14-8. Example of Mesh Group Assignments for Several Ports

6. Repeat step 5 for all ports you want in the mesh domain.

Notes

For meshed ports, leave the **Type** value blank. (Meshed ports do not accept a **Type** value.)

All meshed ports in the switch automatically belong to the same mesh domain. (See figure 14-2 on page 14-4.)

7. When you finish assigning ports to the switch mesh, press [Enter], then [S] (for **Save**). You will then see the following screen.

The asterisk indicates that you must reboot the switch to cause the **Mesh** configuration change to take effect.

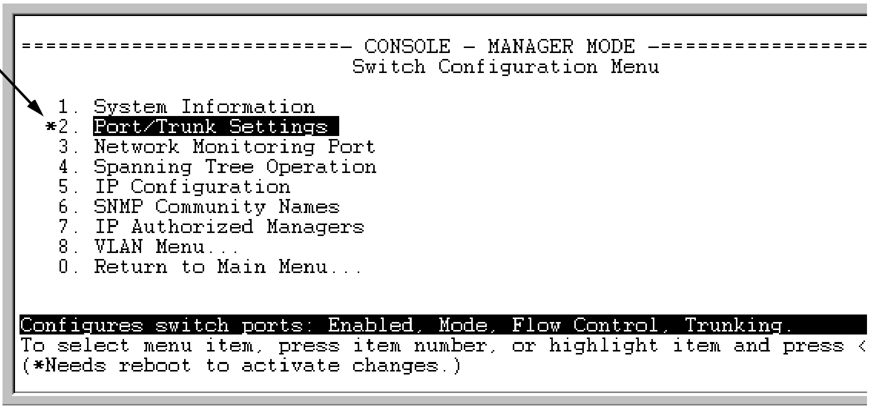

```
========================- CONSOLE - MANAGER MODE -=================
                          Switch Configuration Menu

   1. System Information
  *2. Port/Trunk Settings
   3. Network Monitoring Port
   4. Spanning Tree Operation
   5. IP Configuration
   6. SNMP Community Names
   7. IP Authorized Managers
   8. VLAN Menu...
   0. Return to Main Menu...


Configures switch ports: Enabled, Mode, Flow Control, Trunking.
To select menu item, press item number, or highlight item and press <
(*Needs reboot to activate changes.)
```

Figure 14-9. After Saving a Mesh Configuration Change, Reboot the Switch

8. Press [0] to return to the Main menu.

9. To activate the mesh assignment(s) from the Main menu, reboot the switch by pressing the following keys:

   a. [6] (for **Reboot Switch**)

   b. Space bar (to select **Yes**).

   c. [Enter] (to start the reboot process).

(The switch cannot dynamically reconfigure ports to enable or disable meshing, so it is always necessary to reboot the switch after adding or deleting a port in the switch mesh.)

CLI: To View and Configure Switch Meshing

Port Status and Configuration Features

Feature	Default	Menu	CLI	Web
viewing switch mesh status	n/a	n/a	below	n/a
configuring switch meshing	Disabled	n/a		n/a

Viewing Switch Mesh Status

Syntax: show mesh

Lists the switch ports configured for meshing, along with the **State** *of each mesh-configured connection, the MAC address of the switch on the opposite end of the link (***Adjacent Switch***), and the MAC address of the port on the opposite end of the link (***Peer Port***).*

Reading the Show Mesh Output. For each port configured for meshing, the State column indicates whether the port has an active link to the mesh or is experiencing a problem.

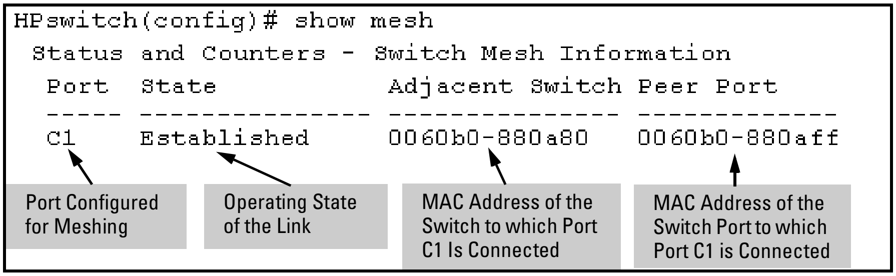

```
HPswitch(config)# show mesh
 Status and Counters - Switch Mesh Information
  Port  State            Adjacent Switch Peer Port
  ----- ---------------- --------------- -------------
  C1    Established      0060b0-880a80   0060b0-880aff
```

Port Configured for Meshing
Operating State of the Link
MAC Address of the Switch to which Port C1 Is Connected
MAC Address of the Switch Port to which Port C1 is Connected

Figure 14-10. Example of the Show Mesh Report

Table 14-1. State Descriptions for Show Mesh Output

State	Meaning
Established	The port is linked to a meshed port on another switch and meshing traffic is flowing across the link. The **show mesh** listing includes the MAC addresses of the adjacent switch and direct connection port on the adjacent switch.
Not Established	The port may be linked to a switch on a port that is not configured for meshing or has gone down.
Initial	The port has just come up as a meshed port and is trying to negotiate meshing.
Disabled	The port is configured for meshing but is not connected to another device.
Error	Indicates a multiple MAC-address error. This occurs when you have two or more mesh ports from the same switch linked together through a hub.
Topology Error	Two meshed switches are connected via a hub and traffic from other, non-meshed devices, is flowing into the hub. The **show mesh** listing includes the MAC addresses of the adjacent switch and direct connection port on the adjacent switch.

Topology Example with Show Mesh. Suppose that you have the following topology:

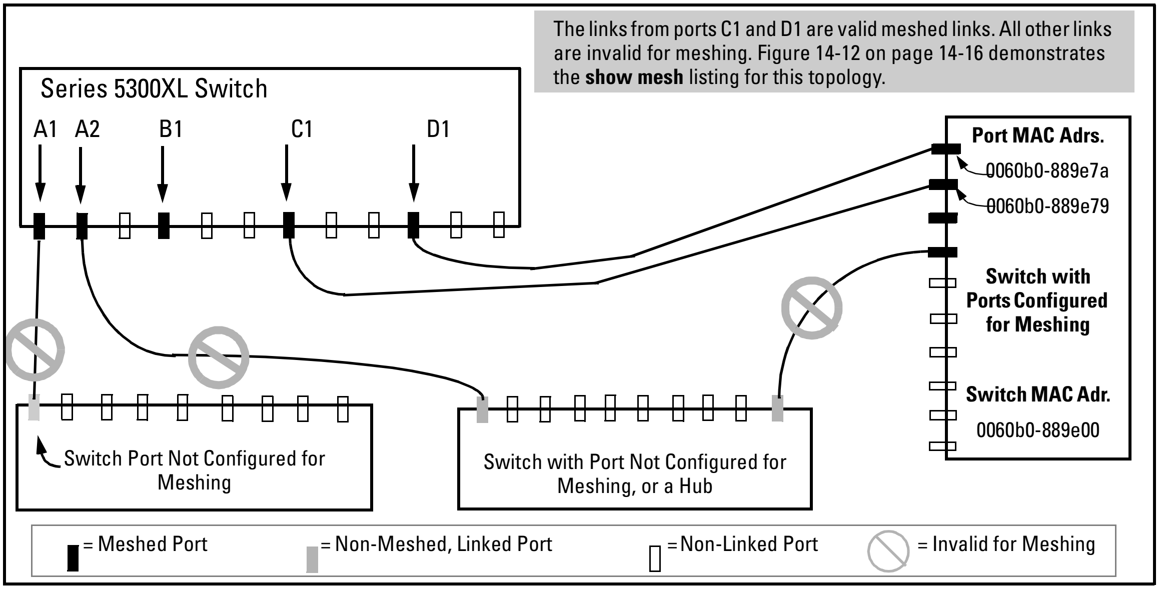

The links from ports C1 and D1 are valid meshed links. All other links are invalid for meshing. Figure 14-12 on page 14-16 demonstrates the **show mesh** listing for this topology.

Series 5300XL Switch
A1
A2
B1
C1
D1
Port MAC Adrs.
0060b0-889e7a
0060b0-889e79
Switch with Ports Configured for Meshing
Switch MAC Adr.
0060b0-889e00
Switch Port Not Configured for Meshing
Switch with Port Not Configured for Meshing, or a Hub
= Meshed Port
= Non-Meshed, Linked Port
= Non-Linked Port
= Invalid for Meshing

Figure 14-11. Example of a Meshed Topology with Some Mesh Ports Incorrectly Linked

Table 14-2 describes the meshing operation in the above topology.

Table 14-2. Operating Details for Figure 14-11

Port	Meshing?	Connection
A1	Yes	Connected to a port that may not be configured for meshing
A2	Yes	Connected to a switch port on a device that is not configured for meshing (another switch, or a hub). In this case, the **Topology Error** message indicates that the switch detects a meshed port on another, non-adjacent device that is also connected to the non-meshed switch or hub. *However, meshing will not operate properly through this connection.*
B1	Yes	Not connected to another device.
C1	Yes	Connected to a meshed port on the same adjacent switch as D1 with meshing operating properly.
D1	Yes	Connected to a meshed port on the same adjacent switch as C1 with meshing operating properly.

Figure 14-12 lists the show mesh display for the topology and meshing configuration in figure 14-11:

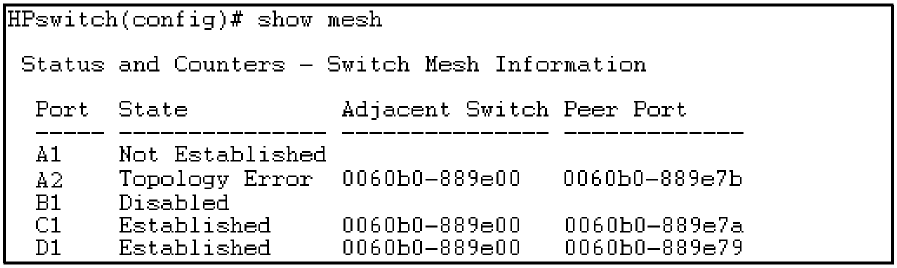

```
HPswitch(config)# show mesh

 Status and Counters - Switch Mesh Information

  Port  State            Adjacent Switch Peer Port
  ----- ---------------- --------------- -------------
  A1    Not Established
  A2    Topology Error   0060b0-889e00   0060b0-889e7b
  B1    Disabled
  C1    Established      0060b0-889e00   0060b0-889e7a
  D1    Established      0060b0-889e00   0060b0-889e79
```

Figure 14-12. Example of the Show Mesh Listing for the Topology in Figure 14-11

CLI: Configuring Switch Meshing

Syntax: [no] mesh [e] < *port-list* >

Enables or disables meshing operation on the specified ports.

All meshed ports on a switch belong to the same mesh domain. Thus, to configure multiple meshed ports on a switch, you need to:

1. Specify the ports you want to operate in the mesh domain.
2. Use **write memory** to save the configuration to the startup-config file.
3. Reboot the switch

For example, to configure meshing on ports A1-A4, B3, C1, and D1-D3:

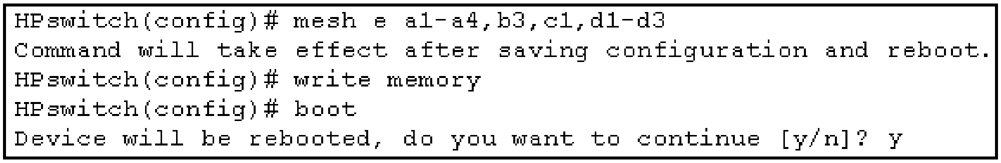

```
HPswitch(config)# mesh e a1-a4,b3,c1,d1-d3
Command will take effect after saving configuration and reboot.
HPswitch(config)# write memory
HPswitch(config)# boot
Device will be rebooted, do you want to continue [y/n]? y
```

Figure 14-13. Example of How To Configure Ports for Meshing

To remove a port from meshing, use the "**no**" version of **mesh**, followed by **write memory** and rebooting the switch. For example, to remove port C1 from the mesh:

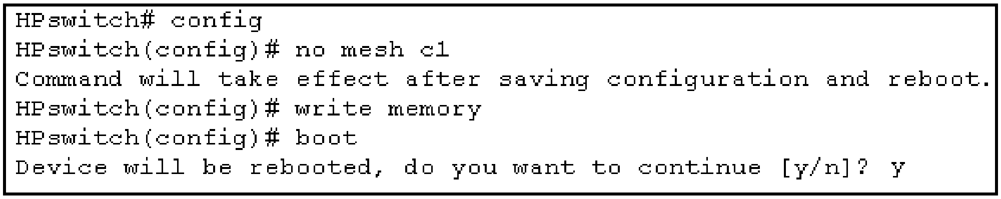

```
HPswitch# config
HPswitch(config)# no mesh c1
Command will take effect after saving configuration and reboot.
HPswitch(config)# write memory
HPswitch(config)# boot
Device will be rebooted, do you want to continue [y/n]? y
```

Figure 14-14. Example of Removing a Port from the Mesh

Operating Notes for Switch Meshing

In a switch mesh domain traffic is distributed across the available paths with an effort to keep latency the same from path to path. The path selected at any time for a connection between a source node and a destination node is based on these latency and throughput cost factors:

- Outbound queue depth, or the current outbound load factor for any given outbound port in a possible path
- Port speed, such as 10Mbps versus 100Mbps; full-duplex or half-duplex
- Inbound queue depth, or how busy is a destination switch in a possible path
- Increased packet drops, indicating an overloaded port or switch

Paths having a lower cost will have more traffic added than those having a higher cost. Alternate paths and cost information is discovered periodically and communicated to the switches in the mesh domain. This information is used to assign traffic paths between devices that are newly active on the mesh.

This means that after an assigned path between two devices has timed out, new traffic between the same two devices may take a different path than previously used.

To display information on the operating states of meshed ports and the identities of adjacent meshed ports and switches, see "Viewing Switch Mesh Status" on page 14-14.

Flooded Traffic

Broadcast and multicast packets will always use the same path between the source and destination edge switches unless link failures create the need to select new paths. (Broadcast and multicast traffic entering the mesh from different edge switches are likely to take different paths.) When an edge switch receives a broadcast from a non-mesh port, it floods the broadcast out all its other non-mesh ports, but sends the broadcast out only those ports in the mesh that represent the path from that edge switch through the mesh domain. (Only one copy of the broadcast packet gets to each edge switch for broadcast out of its non-meshed ports. This helps to keep the latency for these packets to each switch as low as possible.)

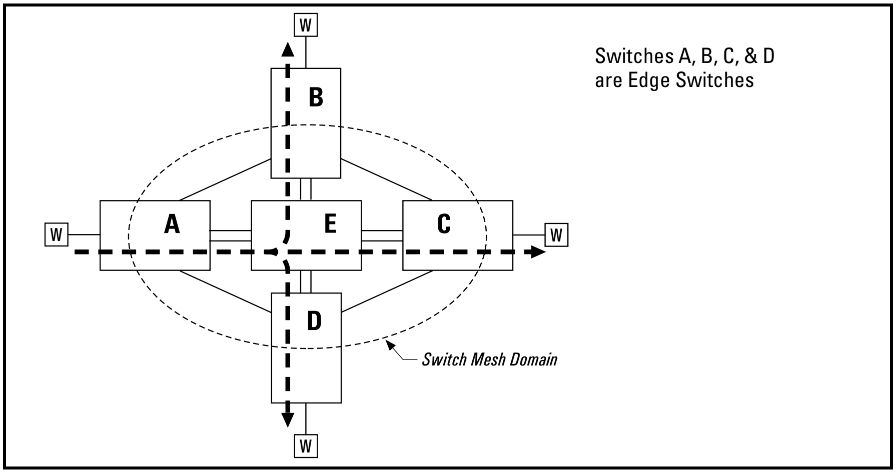

Figure 14-15. Example of a Broadcast Path Through a Switch Mesh Domain

Any mesh switches that are not edge switches will flood the broadcast packets only through ports (paths) that link to separate edge switches in the controlled broadcast tree. The edge switches that receive the broadcast will flood the broadcast out all non-meshed ports. Some variations on broadcast/multicast traffic patterns, including the situation where multiple VLANs are configured and a broadcast path through the mesh domain leads only to ports that are in the same VLAN as the device originating the broadcast.

Unicast Packets with Unknown Destinations

A meshed switch receiving a unicast packet with an unknown destination does not flood the packet onto the mesh. Instead, the switch sends a query on the mesh to learn the location of the unicast destination. The meshed switches then send 802.2 test packets through their non-meshed ports. After the unicast destination is found and learned by the mesh, subsequent packets having the same destination address will be forwarded. By increasing the **MAC Age Time** in the System Information screen (page 6-10) you can cause the switch address table to retain device addresses longer. Because the switches in a mesh exchange address information, this will help to decrease the number of unicast packets with unknown destinations, which improves latency within the switch mesh. Also, in an IP environment, HP recommends that you configure IP addresses on meshed switches. This makes the discovery mechanism more robust, which contributes to decreased latency.

Spanning Tree Operation with Switch Meshing

Using STP or RSTP with several switches and no switch meshing configured can result in unnecessarily blocking links and reducing available bandwidth. For example:

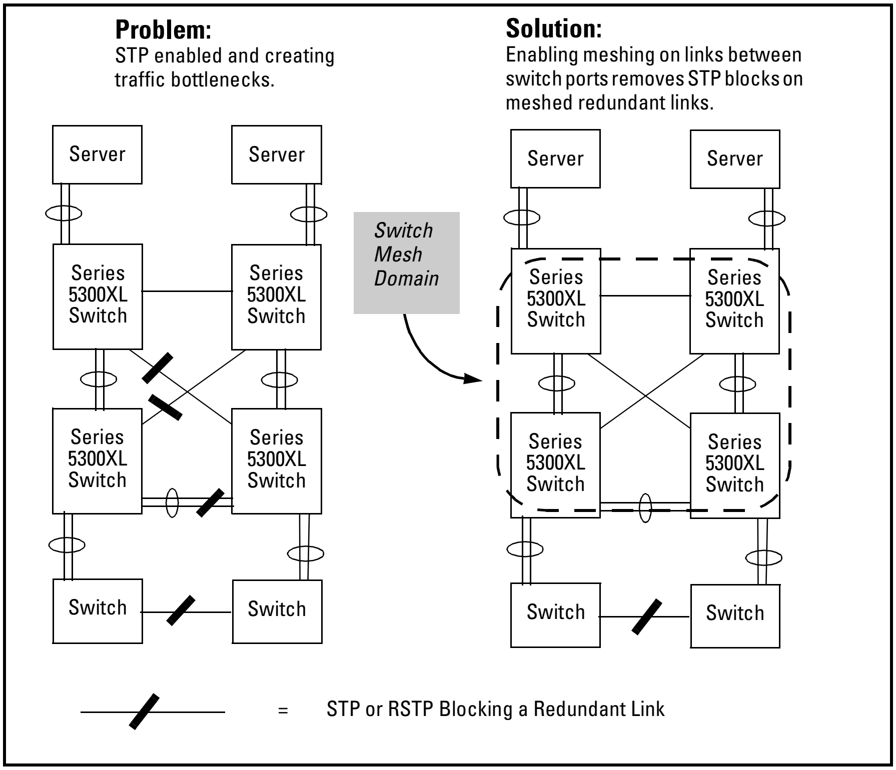

Figure 14-16. Example Using STP Without and With Switch Meshing

If you enable STP or RSTP on any meshed switch, you should enable either STP or RSTP on all switches in the mesh. (That is, if you are going to use spanning-tree in a switch mesh, all switches in the mesh should be configured with the same type of spanning-tree: 802.1d/STP or 802.1w/RSTP.) STP and RSTP see a meshed domain as a single link. However, on edge switches in the domain, STP and RSTP will manage non-meshed redundant links from other devices. For example:

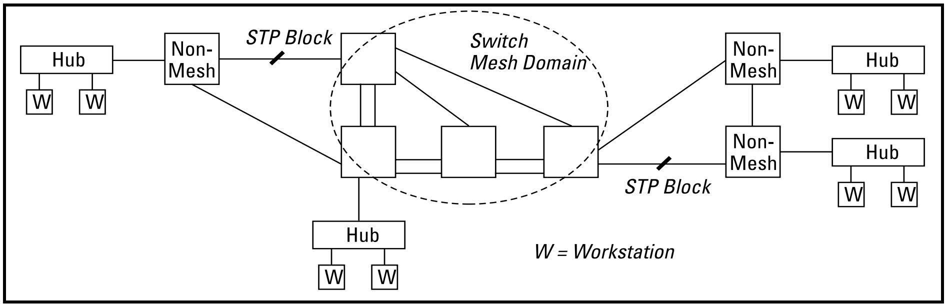

Figure 14-17. Connecting a Switch Mesh Domain to Non-Meshed Devices

STP or RSTP should be configured on non-mesh devices that use redundant links to interconnect with other devices or with multiple switch mesh domains. For example:

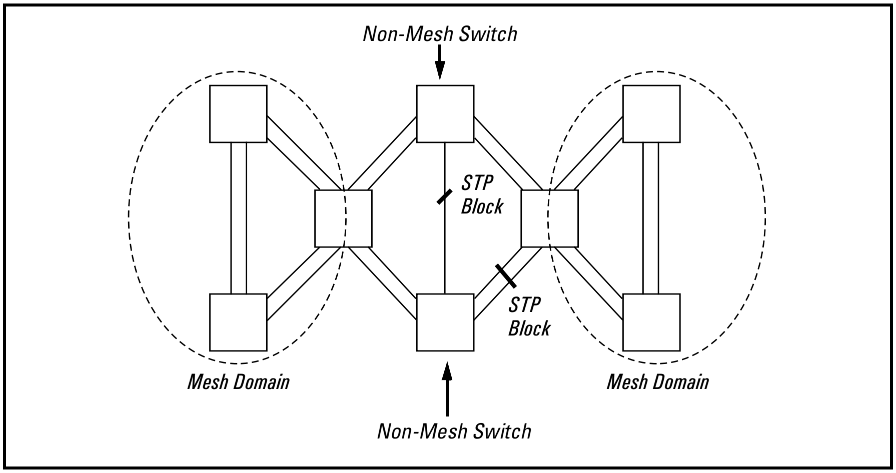

Figure 14-18. Interconnecting Switch Mesh Domains with Redundant Links

In the above case of multiple switch meshes linked with redundant trunks there is the possibility that STP or RSTP will temporarily block a mesh link. This is because it is possible for STP or RSTP to interpret the cost on an external trunked link to be less than the cost on a meshed link. However, if this condition occurs, the meshed switch that has a blocked link will automatically increase the cost on the external (non-meshed) link to the point where STP or RSTP will block the external link and unblock the meshed link. This process typically resolves itself in approximately 30 seconds.

Caution

Because the switch automatically gives faster links a higher priority, the default STP or RSTP parameter settings are usually adequate for spanning tree operation. Because incorrect settings can adversely affect network performance, you should avoid making changes without having a strong understanding of how STP or RSTP operates. For more on STP and RSTP, refer to chapter 13, "802.1w Rapid Spanning Tree Protocol (RSTP) and 802.1d Spanning Tree Protocol". Also, you may want to examine the IEEE 802.1d or 802.1w standard, depending on which version of spanning-tree you are using.

Filtering/Security in Meshed Switches

Because paths through the mesh can vary with network conditions, configuring filters on meshed ports can create traffic problems that are difficult to predict, and is not recommended. However, configuring filters on non-meshed ports in an edge switch provides you with control and predictability.

IP Multicast (IGMP) in Meshed Switches

Like trunked ports, the switch mesh domain appears as a single port to IGMP. However, unlike trunked ports, IGMP protocol and multicast traffic may be sent out over several links in the mesh in the same manner as broadcast packets.

Static VLANs

In a network having a switch mesh domain and multiple static VLANs configured, all static VLANs must be configured on each meshed switch, even if no ports on the switch are assigned to any VLAN. (The switch mesh is a member of all static VLANs configured on the switches in the mesh.)

When static VLANs are configured, the mesh is seen as a single entity by each VLAN. All ports in the mesh domain are members of all VLANs and can be used to forward traffic for any VLAN. However, the non-mesh ports on edge switches that allow traffic to move between the mesh and non-meshed devices belong to specific VLANs and do not allow packets originating in a specific VLAN to enter non-meshed devices that do not belong to that same VLAN. (It is necessary to use a router to communicate between VLANs.) For example, in the following illustration, traffic from host A entering the switch mesh can only exit the mesh at the port for hosts B and E. Traffic from host A for any other host (such as C or D) will be dropped because only hosts B and E are in the same VLAN as host A.

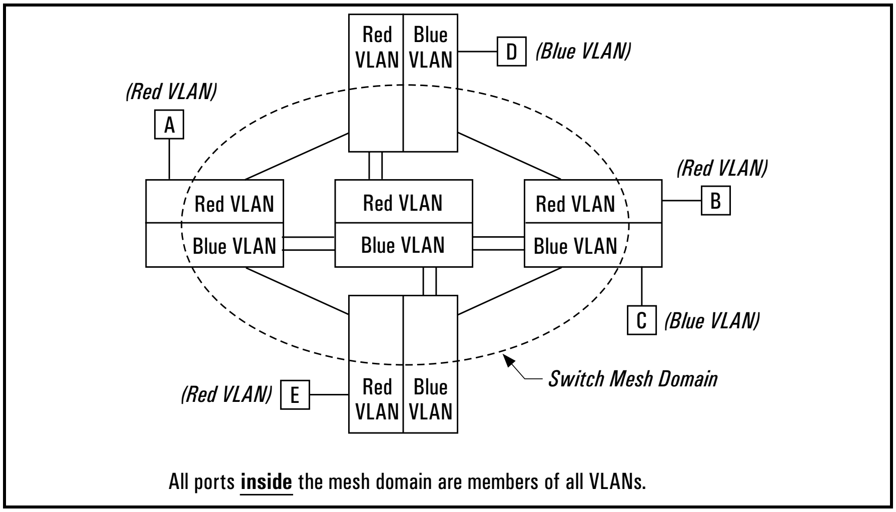

Figure 14-19. VLAN Operation with a Switch Mesh Domain

Dynamic VLANs

If GVRP is enabled, meshed ports in a switch become members of any dynamic VLANs created in the switch in the same way that they would if meshing was not configured in the switch. (For more on GVRP, refer to "GVRP" on page 11-34.)

Requirements and Restrictions

(See also “Operating Rules” on page 14-5.)

- **Mesh Domain Size:** Up to 12 switches are supported in a switch mesh domain. The following example illustrates a meshed backbone where the maximum meshed switch hop count is 3.

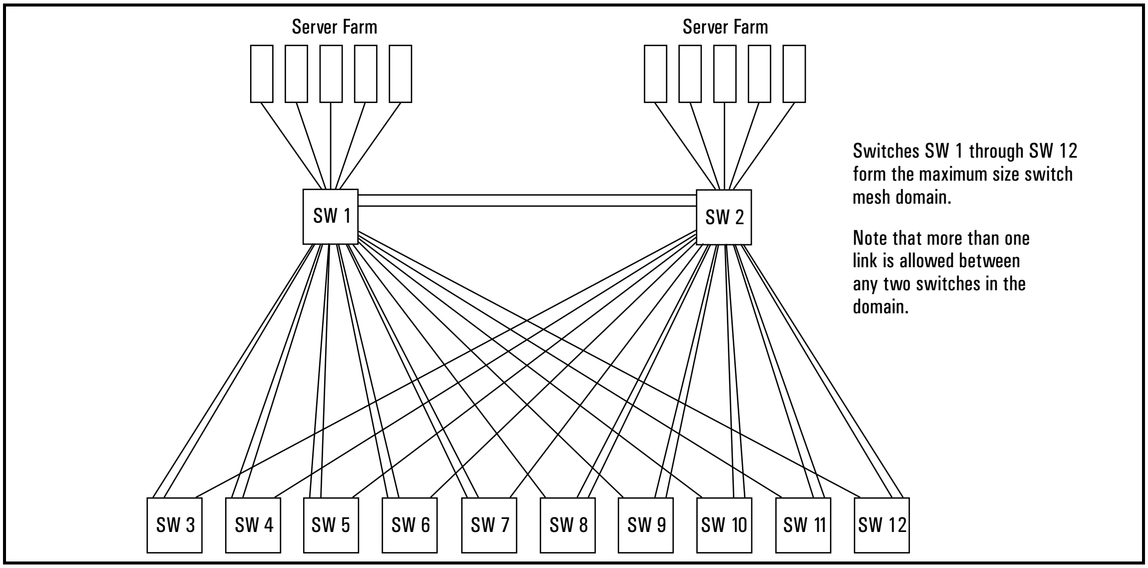

Switches SW 1 through SW 12 form the maximum size switch mesh domain.

Note that more than one link is allowed between any two switches in the domain.

Figure 14-20. Example of a Backbone Using the Maximum Number of Meshed Switches

- **Mesh Support Within the Domain:** All switches in the mesh domain, including edge switches, must support the HP switch meshing protocol.
- **Switch Hop Count in the Mesh Domain:** A maximum (meshed) switch hop count of five is allowed in the path connecting two nodes via a switch mesh domain topology.

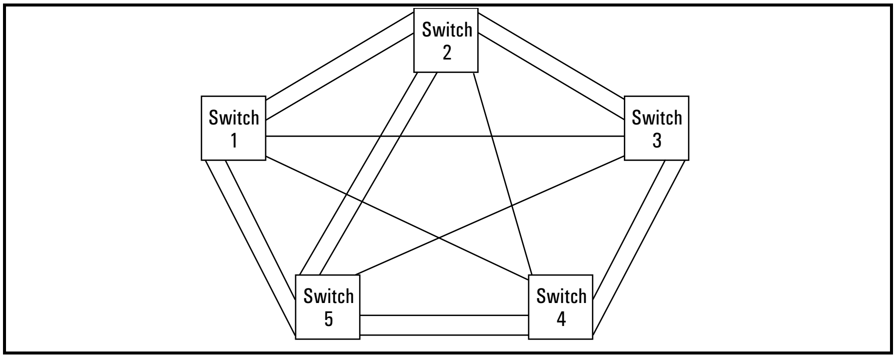

Figure 14-21. Example of the Maximum Meshed Switch Hop Count

- **Connecting Mesh Domains:** To connect two separate switch meshing domains, you must use non-meshed ports. (The non-meshed link can be a port trunk or a single link.) Refer to figure 14-3 on page 14-7.
- **Fast EtherChannel® (FEC):** This cannot be configured on a meshed port. (You can configure FEC on non-meshed ports in a switch that also has meshed ports.)
- **Multiple Links Between Meshed Switches:** Multiple mesh ports can be connected between the same two switches, to provide higher bandwidth. Each port that you want in the mesh domain should be configured as **Mesh** (and not as a trunk—**Trk**). Note that if you configure a port as **Mesh**, there is no "Type" selection for that port.
- **Automatic Broadcast Control:** Series 5300XL switches do not offer this feature. Thus, in a switch mesh comprised of Series 5300XL switches and any of the 1600M/2400M/2424M/4000M/8000M switches, ABC must be disabled (the default setting) on the 1600M/2400M/2424M/4000M/8000M switches.
- **Network Monitor Port:** If a network monitor port is configured, broadcast packets may be duplicated on this port if more than one port is being monitored and switch meshing is enabled.
- **Compatibility with Older Switches:** Each entry in a Series 5300XL switch MAC-address table consists of a MAC address and a VLAN ID (VID). In older switches there is no VID; just a MAC address. The older switches will therefore see indistinguishable, duplicate addresses where the Series 5300XL switches will see multiple different addresses consisting on the same MAC address and different VIDs. In a switch mesh that includes any 1600M/2400M/2424M/4000M/8000M switches, duplicate MAC addresses entering the mesh on different switches are not allowed. (These older switches do not recognize multiple instances of a particular MAC address on different VLANs.) If you try to add one of these switches to a mesh comprised entirely of Series 5300XL switches, and any of the

Series 5300XL switches detects a duplicate MAC address entering the mesh through separate switches, the 1600M/2400M/2424M/4000M/8000M switch will not be allowed in the switch mesh.

For additional information on troubleshooting meshing problems, refer to "Using a Heterogeneous Switch Mesh" on page 14-7 and "Mesh-Related Problems" on page C-9.

Quality of Service (QoS): Managing Bandwidth More Effectively

Contents

Introduction

QoS Feature	Default	Menu	CLI	Web
UDP/TCP Priority	Disabled	—	page 15-14	page 15-57
IP-Device Priority	Disabled	—	page 15-20	"
IP Type-of-Service Priority	Disabled	—	page 15-25	"
LAN Protocol Priority	Disabled	—	page 15-37	"
VLAN-ID Priority	Disabled	—	page 15-39	"
Source-Port Priority	Disabled	—	page 15-44	"
DSCP Policy Table	Various	—	page 15-51	"

As the term suggests, *network policy* refers to the network-wide controls you can implement to:

- Ensure uniform and efficient traffic handling throughout your network, while keeping the most important traffic moving at an acceptable speed, regardless of current bandwidth usage.
- Exercise control over the priority settings of inbound traffic arriving in and travelling through your network.

Adding bandwidth is always a good idea, but it is not always feasible and does not completely eliminate the potential for network congestion. There will always be points in the network where multiple traffic streams merge or where network links will change speed and capacity. The impact and number of these congestion points will increase over time as more applications and devices are added to the network.

When (not *if*) network congestion occurs, it is important to move traffic on the basis of relative importance. However, without *Quality of Service* (QoS) prioritization, less important traffic can consume network bandwidth and slow down or halt the delivery of more important traffic. That is, without QoS, most traffic received by the switch is forwarded with the same priority it had upon entering the switch. In many cases, such traffic is "normal" priority and competes for bandwidth with all other normal-priority traffic, regardless of its relative importance to your organization's mission.

This section gives an overview of QoS operation and benefits, and describes how to configure QoS in the console interface.

Quality of Service is a general term for classifying and prioritizing traffic throughout a network. That is, QoS enables you to establish an end-to-end traffic priority policy to improve control and throughput of important data. You can manage available bandwidth so that the most important traffic goes first. For example, you can use Quality of Service to:

- Upgrade or downgrade traffic from various servers.
- Control the priority of traffic from dedicated VLANs or applications.
- Change the priorities of traffic from various segments of your network as your business needs change.
- Set priority policies in edge switches in your network to enable traffic-handling rules across the network.

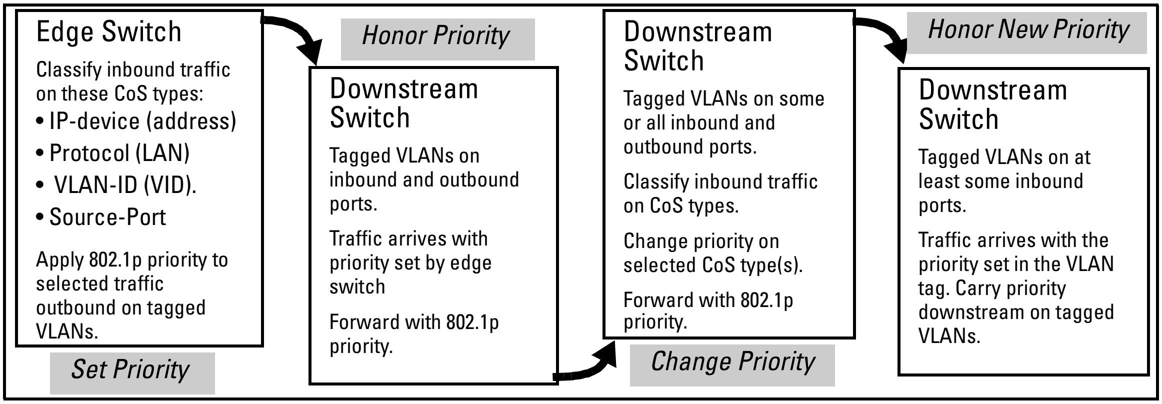
Edge Switch
Classify inbound traffic on these CoS types:
• IP-device (address)
• Protocol (LAN)
• VLAN-ID (VID).
• Source-Port
Apply 802.1p priority to selected traffic outbound on tagged VLANs.
Set Priority
Honor Priority
Downstream Switch
Tagged VLANs on inbound and outbound ports.
Traffic arrives with priority set by edge switch
Forward with 802.1p priority.
Downstream Switch
Tagged VLANs on some or all inbound and outbound ports.
Classify inbound traffic on CoS types.
Change priority on selected CoS type(s).
Forward with 802.1p priority.
Change Priority
Honor New Priority
Downstream Switch
Tagged VLANs on at least some inbound ports.
Traffic arrives with the priority set in the VLAN tag. Carry priority downstream on tagged VLANs.

Figure 15-1. Example of 802.1p Prioritization Based on CoS Types and Use of Tagged VLANs

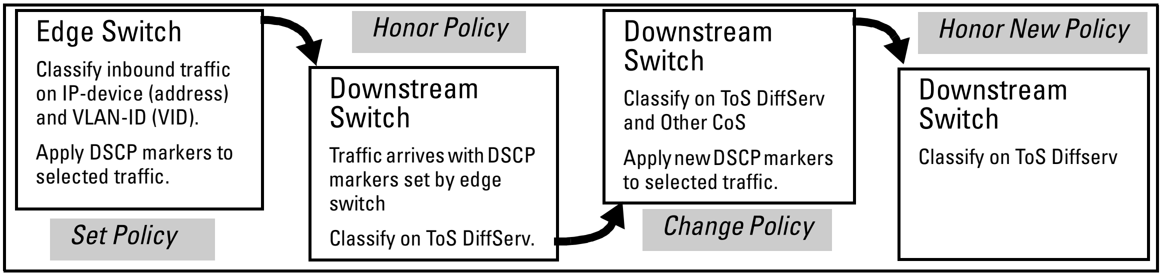
Edge Switch
Classify inbound traffic on IP-device (address) and VLAN-ID (VID).
Apply DSCP markers to selected traffic.
Set Policy
Honor Policy
Downstream Switch
Traffic arrives with DSCP markers set by edge switch
Classify on ToS DiffServ.
Downstream Switch
Classify on ToS DiffServ and Other CoS
Apply new DSCP markers to selected traffic.
Change Policy
Honor New Policy
Downstream Switch
Classify on ToS Diffserv

Figure 15-2. Example Application of Differentiated Services Codepoint (DSCP) Policies

At the edge switch, QoS classifies certain traffic types and in some cases applies a DSCP policy. At the next hop (downstream switch) QoS honors the policies established at the edge switch. Further downstream, another switch may reclassify some traffic by applying new policies, and yet other downstream switches can be configured to honor the new policies.

QoS is implemented in the form of rules or policies that are configured on the switch. While you can use QoS to prioritize only the outbound traffic moving through the switch, you derive the maximum benefit by using QoS in an 802.1Q VLAN environment (with 802.1p priority tags) or in an untagged VLAN environment (with DSCP policies) where QoS can set priorities that downstream devices can support without re-classifying the traffic.

By prioritizing traffic, QoS supports traffic growth on the network while optimizing the use of existing resources—and delaying the need for further investments in equipment and services. That is, QoS enables you to:

- Specify which traffic has higher or lower priority, regardless of current network bandwidth or the relative priority setting of the traffic when it is received on the switch.
- Change (upgrade or downgrade) the priority of outbound traffic.
- Override "illegal" packet priorities set by upstream devices or applications that use 802.1Q VLAN tagging with 802.1p priority tags.
- Avoid or delay the need to add higher-cost NICs (network interface cards) to implement prioritizing. (Instead, control priority through network policy.)

QoS on the Series 5300XL switches supports these types of traffic marking:

- **802.1p prioritization:** Controls the outbound port queue priority for traffic leaving the switch, and (if traffic exits through a tagged VLAN) sends the priority setting with the individual packets to the downstream devices.
- **IP Type-of-Service (ToS):** Enables the switch to set, change, and honor prioritization policies by using the Differentiated Services (diffserv) bits in the ToS byte of IPv4 packet headers.

Terminology

Term	Use in This Document
codepoint	Refer to DSCP, below.
downstream device	A device linked directly or indirectly to an outbound switch port. That is, the switch sends traffic to downstream devices.
DSCP	*Differentiated Services Codepoint.* (Also termed *codepoint.*) A DSCP is comprised of the upper six bits of the ToS (Type-of-Service) byte in IP packets. There are 64 possible codepoints. In the switch's default QoS configuration, some codepoints are configured with default 802.1p priority settings for Assured-Forwarding and Expedited Forwarding, while others are unused (and listed with **No-override** for a priority).
DSCP policy	A DSCP configured with a specific 802.1p priority (0- 7). (Default: **No-override**). Using a DSCP policy you can configure the switch to assign priority to IP packets. That is, for an IP packet identified by the specified classifier, you can assign a new DSCP and an 802.1p priority (0-7). For more on DSCP, refer to "Details of QoS IP Type-of-Service" on page 15-34. For the DSCP map, see figure 15-18 on page 15-35.
edge switch	In the QoS context, this is a switch that receives traffic from outside the LAN and forwards it to devices within the LAN. Typically, an edge switch is used with QoS to recognize packets based on classifiers such as TCP/UDP application type, IP-device (address), Protocol (LAN), VLAN-ID (VID), and Source-Port (although it can also be used to recognize packets on the basis of ToS bits). Using this packet recognition, the edge switch can be used to set 802.1p priorities or DSCP policies that downstream devices will then honor.
inbound port	Any port on the switch through which traffic enters the switch.
IPv4	Version 4 of the IP protocol.
outbound packet	A packet leaving the switch through any LAN port.
outbound port	Any port on the switch through which traffic leaves the switch.
outbound port queue	For any port, a buffer that holds outbound traffic until it can leave the switch through that port. There four outbound queues for each port in the switch: high, medium, normal, and low. Traffic in a port's high priority queue leaves the switch before any traffic in the port's medium priority queue, and so-on.
IP-precedence bits	The upper three bits in the Type of Service (ToS) field of an IP packet.
upstream device	A device linked directly or indirectly to an inbound switch port. That is, the switch receives traffic from upstream devices.
802.1p priority	A traffic priority setting carried by packets moving from one device to another in an 802.1Q tagged VLAN environment. This setting can be from 0 - 7. The switch handles an outbound packet on the basis of its 802.1p priority. However, if the packet leaves the switch through an untagged VLAN, this priority is dropped, and the packet arrives at the next, downstream device without an 802.1p priority assignment.
802.1Q tagged VLAN	A virtual LAN (VLAN) that complies with the 802.1Q standard and is configured as "tagged". In this environment, IP packets carry an 802.1p priority from one device to the next.

Overview

QoS settings operate on two levels:

- **Controlling the priority of outbound packets moving through the switch:** Each switch port has four outbound traffic queues; “low”, “normal”, “medium”, and “high” priority. Packets leave the switch port on the basis of their queue assignment and whether any higher queues are empty:

  Table 15-1.Port Queue Exit Priorities

Port Queue	Priority for Exiting From the Port
High (6 - 7)	First
Medium (4 - 5)	Second
Normal (0, 3)	Third
Low (1 - 2)	Fourth

  A QoS configuration enables you to set the outbound priority queue to which a packet is sent. (In an 802.1Q tagged VLAN environment, if QoS is *not* configured on the switch, but *is* configured on an upstream device, the priorities carried in the packets determine the forwarding queues in the switch.)

- **Configuring a priority for outbound packets and a service (priority) policy for use by downstream devices:**
  - **DSCP Policy:** This feature enables you to set a priority policy in outbound IP packets. (You can configure downstream devices to read and use this policy.) This method is not dependent on tagged VLANs to carry priority policy to downstream devices, and can:
    – Change the codepoint (the upper six bits) in the TOS byte.
    – Set a new 802.1p priority for the packet.

    (Setting DSCP policies requires IPv4 inbound packets. Refer to the “IPv4” entry under “Terminology” on page 15-5.)
  - **802.1p Priority:** If an outbound packet is in an 802.1Q tagged VLAN environment (that is, if the packet is assigned to a tagged VLAN on the outbound port), then the packet carries an 802.1p priority setting that was configured in the switch. This priority setting ranges from 0 to 7, and can be used by downstream devices having up to eight outbound port queues. Thus, while packets within the switch move at the four priority levels shown in table 15-1, above, they still can carry an 802.1p priority that can be used by downstream devices having more or less than the four priority levels in the Series 5300XL switches. Also, if the packet enters the switch with an 802.1p priority setting, QoS can override this setting if configured to do so.

Note: If you are not using multiple tagged VLANs in your network, you can still use the tagged VLAN feature by configuring the default VLAN as a tagged VLAN.

You can configure a QoS priority of 0 through 7 for an outbound packet. When the packet is then sent to a port, the QoS priority determines which outbound queue the packet uses:

Table 15-2. QoS Priority Settings and Operation

QoS Priority Setting	Outbound Port Queue
1 - 2	low priority
0 - 3	normal priority
4 - 5	medium priority
6 - 7	high priority

If a packet is not in an 802.1Q tagged VLAN environment, then the above QoS settings (table 15-2) control only to which outbound queue the packet goes, and no 802.1p priority is added to the packet for use by downstream devices. But if the packet is in an 802.1Q tagged VLAN environment, then the above setting is also added to the packet as an 802.1p priority that can be used by downstream devices and applications, as indicated in the next table. In either case, an IP packet can also carry a prioritization policy to downstream devices by using codepoint-marking in the ToS byte.

Table 15-3. Mapping Series 5300XL QoS Priority Settings to Device Queues

Priority Setting in the Series 5300XLSwitches	Outbound Port Queues in the Series 5300XL Switches	802.1p Priority Setting Added to Tagged VLAN Packets Leaving the Switch	Queue Assignment in Downstream Devices With:		
			8 Queues	3 Queues	2 Queues
1	Queue 1	1 (low priority)	Queue 1	Queue 1	Queue 1
2		2	Queue 2		
0	Queue 2	0 (normal priority)	Queue 3	Queue 2	
3		3	Queue 4		
4	Queue 3	4 (medium priority)	Queue 5	Queue 3	Queue 2
5		5	Queue 6		
6	Queue 4	6 (high priority)	Queue 7		
7		7	Queue 8		

Classifiers for Prioritizing Outbound Packets

You can configure QoS prioritization on the basis of seven *QoS classifiers*, or packet criteria, evaluated in the following order:

Precedence	QoS Classifier
1	UDP/TCP Application Type (port)
2	Device Priority (destination or source IP address)
3	IP Type of Service (ToS) field (IP packets only)
4	Protocol Priority (IP, IPX, ARP, DEC LAT, AppleTalk, SNA, and NetBeui)
5	VLAN Priority
6	Incoming source-port on the switch
7	Incoming 802.1p Priority (present in tagged VLAN environments)

If the switch is configured with multiple classifiers that address the same packet, the switch uses only the QoS configuration for the QoS classifier that has the highest precedence. (In this case, the QoS configuration for another, lower-precedence classifier that may apply is ignored.) For example, if QoS assigns high priority to “red” VLAN packets, but normal priority to IP packets, since Protocol Priority (4) has precedence over VLAN priority (5), IP packets on the “red” VLAN will be set to normal priority. See Table 15-4. "Precedence Criteria for QoS Classifiers" on page 15-9 for more information.

Note On Using Multiple Criteria

HP recommends that you configure a minimum number of the available QoS classifiers for prioritizing any given packet type. Increasing the number of active classifier options for a packet type increases the complexity of the possible outcomes.

Table 15-4. Precedence Criteria for QoS Classifiers

Precedence	Criteria	Overview
1	UDP/TCP	Takes precedence based on a layer 4 UDP or TCP application, with a user-specified application port number (for example, Telnet). **Default state:** Disabled
	If a packet does not meet the criteria for UDP/TCP priority, then precedence defaults to the Device Priority classifier, below.	
2	Device Priority (IP Address)	Takes precedence based on an outbound packet having a particular destination or source IP address. QoS allows up to 256 IP addresses. If an outbound packet has an IP address as the destination, it takes precedence over another outbound packet that has the same IP address as a source. (This can occur, for example, on an outbound port in a switch mesh environment.) Also, if the source and destination IP addresses (SA and DA) in the same packet match, the DA takes precedence. **Default state:** No IP address prioritization.
	If a packet does not meet the criteria for device priority, then precedence defaults to the IP Type of Service (ToS) classifier, below.	
3	IP Type-of-Service (ToS)	Takes precedence based on the TOS field in IP packets. (Applies only to IP packets.) The ToS field is configured by an upstream device or application before the packet enters the switch. • **IP Precedence Mode:** QoS reads the packet's IP precedence (upper three) bits in the Type-of-Service (ToS) byte and automatically prioritizes the packet (if specified in the QoS configuration) for outbound transmission. • **Differentiated Services Mode:** QoS reads the packet's differentiated services, or codepoint (upper six) bits of the Type-of-Service (TOS) byte. Packet prioritization depends on the configured priority for the codepoint. (Some codepoints default to the DSCP standard, but can be overridden.) For more on this topic, see "QoS IP Type-of-Service (ToS) Policy and Priority" on page 15-25. **Default state:** Disabled.
	If a packet does not meet the criteria for ToS priority, then precedence defaults to the Protocol classifier, below.	
4	Layer 3 Protocol Priority	Takes precedence based on network protocols: IP, IPX, ARP, DEC LAT, AppleTalk, SNA, and NetBeui. **Default state:** No-override for any protocol.
	If a packet does not meet the criteria for Protocol priority, then precedence defaults to the VLAN classifier, below.	
5	VLAN Priority	Takes precedence based on the ID number of the VLAN in which the packet exists. For example, if the default VLAN (VID = 1) and the "Blue" VLAN (with a VID of 20) are both assigned to a port, and Blue VLAN traffic is more important, you can configure QoS to give Blue VLAN traffic a higher priority than default VLAN traffic. (Priority is applied on the outbound port.) **Default state:** No-override.
	If a packet does not meet the criteria for VLAN priority, then precedence defaults to the Source-Port classifier, below.	

<table>
<tr><th>Precedence</th><th>Criteria</th><th>Overview</th></tr>
<tr><td>6</td><td>Source-Port</td><td>Takes precedence based on the source-port (that is, the port on which the packet entered the switch).</td></tr>
<tr><td></td><td colspan="2">If a packet does not meet the criteria for source-port priority, then precedence defaults to Incoming 802.1p criteria, below</td></tr>
<tr><td>7</td><td>Incoming 802.1p Priority</td><td>Where a packet enters the switch on a tagged VLAN, if QoS is not configured to override the packet's priority setting, the switch uses the packet's existing 802.1p priority (assigned by an upstream device or application) to determine which inbound and outbound port queue to use. If the packet leaves the switch on a tagged VLAN, then there is no change to its 802.1p priority setting. If the packet leaves the switch on an untagged VLAN, the 802.1p priority is dropped.
<table>
<tr><th>Entering (Inbound) 802.1p Priority</th><th>Outbound Port Queue</th><th>Exiting (Outbound) 802.1p Priority</th></tr>
<tr><td>1 - 2</td><td>Low</td><td>1 - 2</td></tr>
<tr><td>0 - 3</td><td>Normal</td><td>0 - 3</td></tr>
<tr><td>4 - 5</td><td>Medium</td><td>4 - 5</td></tr>
<tr><td>6 - 7</td><td>High</td><td>6 - 7</td></tr>
</table>
</td></tr>
<tr><td></td><td colspan="2">If a packet does not meet the criteria for Incoming 802.1p priority, then the packet goes to the "normal" outbound queue of the appropriate port. If the packet entered the switch on an untagged VLAN, but exits on a tagged VLAN, then a tagged VLAN field, including an 802.1p priority of 0 (normal), is added to the packet.</td></tr>
</table>

No Override. By default, the IP ToS, Protocol, and VLAN-ID criteria automatically list each of their priority options as **No-override**. (Some IP TOS codepoints use default priority settings defined by the DSCP standard.) This means that if you do not configure a priority for a specific option, QoS does not prioritize packets to which that option applies. For example, if you do not specify a priority for the IP protocol, then the IP protocol will not be a criteria for setting a QoS priority and the packets will be handled as described above.

Preparation for Configuring QoS

You can use QoS regardless of whether your network has tagged VLANs:

Table 15-5. Summary of QoS Capabilities

Outbound Packet Options	Tagged VLAN Environment	No Tagged VLANs
Control Queue Priority for Packet Types	Yes	Yes
Carry the 802.1p Priority Assignment to Next Downstream Device	Yes	No
Configure a Service Policy and Carry It to Downstream Devices. The policy includes:	Yes [1]	Yes [1]
Assigning a ToS Codepoint		
Assigning an 802.1p Priority [2] to the Codepoint		

[1] Except for packets processed using either the (Layer 3) Protocol or QoS IP-Precedence methods. Also, to use a service policy in this manner, the downstream devices must be configured to interpret and use the DSCP carried in the IP packets.

[2] This priority corresponds to the 802.1p priority scheme and is used to determine the packet's port queue priority. When used in a tagged VLAN environment, this priority is also assigned as the 802.1p priority carried outbound in tagged VLAN packets.

Steps for Configuring QoS on the Switch

1. Determine the QoS policy you want to implement. This includes analyzing the types of traffic flowing through your network and identifying one or more traffic types to prioritize. In order of precedence, these are:

   a. UDP/TCP applications

   b. Device Priority—destination or source IP address (Note that destination has precedence over source. See Table 15-4.)

   c. IP Type-of-Service Precedence Bits (Leftmost three bits in the ToS field of IP packets)

   d. IP Type-of-Service Differentiated Service bits (Leftmost six bits in the ToS field of IP packets)

   e. Protocol Priority

   f. VLAN Priority (requires at least one tagged VLAN on the network)

   g. Source-Port

   h. Incoming 802.1p Priority (requires at least one tagged VLAN on the network)

   For more on how QoS operates with the above traffic types, see Table 15-4. "Precedence Criteria for QoS Classifiers" on page 15-9.)

2. Select the QoS option you want to use. Table 15-6 lists the traffic types (QoS classifiers) and the QoS options you can use for prioritizing or setting a policy on these traffic types:

Table 15-6. Applying QoS Options to Traffic Types Defined by QoS Classifiers

QoS Options for Prioritizing Outbound Traffic		QoS Classifiers						
		UDP/ TCP	IP Device	IP-ToS Precedence	IP- DiffServ	L3 Protocol	VLAN -ID	Source -Port
Option 1: Set 802.1p Priority Only	Prioritize traffic by sending specific packet types (determined by QoS classifier) to different outbound port queues on the switch. Rely on tagged VLANs to carry packet priority as an 802.1p value to downstream devices.	Yes	Yes	Yes [1]	Yes	Yes	Yes	Yes
Option 2: Configure ToS DSCP Policy with 802.1p Priority	Prioritize traffic by sending specific packet types (determined by QoS classifier) to different outbound port queues on the switch. Propagate a service policy by reconfiguring the DSCP in outbound IP packets according to packet type. (Assumes that downstream devices can be configured to recognize the DSCP in IP packets and implement the service policy it indicates.) Use tagged VLANs to carry packet priority as an 802.1p value to downstream devices.	Yes	Yes	No	Yes	No	Yes	Yes

[1] In this mode the configuration is fixed. You cannot change the automatic priority assignment when using IP-ToS Precedence as a QoS classifier.

3. If you want 802.1p priority settings to be included in outbound packets, ensure that tagged VLANs are configured on the appropriate links.

4. Determine the actual QoS configuration changes you will need to make on each QoS-capable device in your network in order to implement the desired policy. Also, if you want downstream devices to read and use DSCPs in IP packets from the switch, configure them to do so by enabling ToS Differentiated Service mode and making sure the same DSCP policies are configured.

Note

If you use TopTools for Hubs & Switches to configure QoS policy in a network, it overrides any QoS settings configured through the CLI or the web browser interface in any individual HP switch.

Using QoS Classifiers To Configure Quality of Service for Outbound Traffic

QoS Feature	Default	Menu	CLI	Web
UDP/TCP Priority	Disabled	—	page 15-14	page 15-57
IP-Device Priority	Disabled	—	page 15-20	“
IP Type-of-Service Priority	Disabled	—	page 15-25	“
LAN Protocol Priority	Disabled	—	page 15-37	“
VLAN-ID Priority	Disabled	—	page 15-39	“
Source-Port Priority	Disabled	—	page 15-44	“

QoS UDP/TCP Priority

QoS Classifier Precedence: 1

When you use UDP or TCP and a layer 4 Application port number as a QoS classifier, traffic carrying the specified UDP/TCP port number(s) is marked with the UDP/TCP classifier's configured priority level, without regard for any other QoS classifiers in the switch.

Options for Assigning Priority. Priority control options for TCP or UDP packets carrying a specified TCP or UDP port number include:

- 802.1p priority
- DSCP policy (Assigning a new DSCP and an associated 802.1p priority; inbound packets must be IPv4.)

For a given TCP or UDP port number, you can use only one of the above options at a time. However, for different port numbers, you can use different options.

TCP/UDP Port Number Ranges. There are three ranges:

- Well-Known Ports: 0 - 1023
- Registered Ports: 1024 - 49151
- Dynamic and/or Private Ports: 49152 - 65535

For more information, including a listing of UDP/TCP port numbers, go to the *Internet Assigned Numbers Authority* (IANA) website at:

http://www.iana.org

Then click on:

Protocol Number Assignment Services

P (Under "Directory of General Assigned Numbers" heading)

Port Numbers

Assigning a Priority Based on TCP or UDP Port Number

This option assigns an 802.1p priority to outbound TCP or UDP packets as described below.

Syntax: qos < udp-port | tcp-port > < *tcp or udp port number* > priority < 0 - 7 >

Configures an 802.1p priority for outbound packets having the specified TCP or UDP application port number. This priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. (Default: Disabled)

no qos < udp-port | tcp-port > < t*cp-udp port number* >

Deletes the specified UDP or TCP port number as a QoS classifier.

show qos tcp-udp-port-priority

Displays a listing of all TCP and UDP QoS classifiers currently in the running-config file.

For example, configure and list 802.1p priority for the following UDP and TCP port prioritization:

TCP/UDP Port	802.1p Priority for TCP	802.1p Priority for UDP
TCP Port 23 (Telnet)	7	7
UDP Port 23 (Telnet)	7	7
TCP Port 80 (World Wide Web HTTP)	2	2
UDP Port 80 (World Wide Web HTTP)	1	1

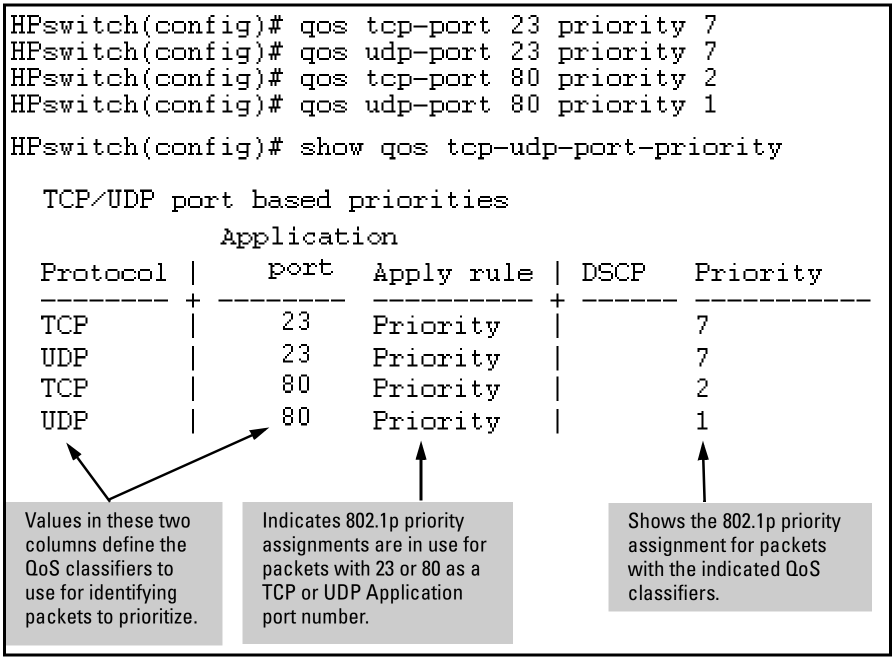

```
HPswitch(config)# qos tcp-port 23 priority 7
HPswitch(config)# qos udp-port 23 priority 7
HPswitch(config)# qos tcp-port 80 priority 2
HPswitch(config)# qos udp-port 80 priority 1

HPswitch(config)# show qos tcp-udp-port-priority

 TCP/UDP port based priorities

              Application
 Protocol |     port     Apply rule | DSCP   Priority
 -------- + --------  ---------- + ------ -----------
 TCP      |     23       Priority   |          7
 UDP      |     23       Priority   |          7
 TCP      |     80       Priority   |          2
 UDP      |     80       Priority   |          1
```

Values in these two columns define the QoS classifiers to use for identifying packets to prioritize.

Indicates 802.1p priority assignments are in use for packets with 23 or 80 as a TCP or UDP Application port number.

Shows the 802.1p priority assignment for packets with the indicated QoS classifiers.

Figure 15-3. Example of Configuring and Listing 802.1p Priority Assignments on TCP/UDP Ports

Assigning a DSCP Policy Based on TCP or UDP Port Number

This option assigns a previously configured DSCP policy (codepoint and 802.1p priority) to outbound TCP or UDP packets having the specified port number. That is, the switch:

1. Selects an incoming IP packet if the TCP or UDP port number it carries matches the port number specified in the TCP or UDP classifier (as shown in figure 15-3, above).
2. Overwrites the packet's DSCP with the DSCP configured in the switch for such packets.
3. Assigns the 802.1p priority configured in the switch for the new DSCP. (Refer to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)
4. Forwards the packet through the appropriate outbound port queue.

For more on DSCP, refer to "Terminology" on page 15-5.

Steps for Creating a DSCP Policy Based on TCP/UDP Port Number Classifiers. This procedure creates a DSCP policy for IP packets carrying the selected UDP or TCP port-number classifier.

1. Identify the TCP or UDP port-number classifier you want to use for assigning a DSCP policy.
2. Determine the DSCP policy for packets carrying the selected TCP or UDP port number.
   a. Determine the DSCP you want to assign to the selected packets. (This codepoint will be used to overwrite the DSCP carried in packets received from upstream devices.)
   b. Determine the 802.1p priority you want to assign to the DSCP.
3. Configure the DSCP policy by using **qos dscp-map** to configure the priority to the codepoint you selected in step 2a. (For details, refer to the example later in this section, and to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)

Note

A codepoint must have an 802.1p priority assignment (0 - 7) before you can configure a policy for prioritizing packets by TCP or UDP port numbers. If a codepoint you want to use shows **No-override** in the **Priority** column of the DSCP map (**show qos dscp-map**), then you must assign a 0 - 7 priority before proceeding.

4. Configure the switch to assign the DSCP policy to packets with the specified TCP or UDP port number.

Syntax: qos dscp-map < *codepoint*> priority < 0 - 7 >

This command is optional if a priority has already been assigned to the < codepoint>. *The command creates a DSCP policy by assigning an 802.1p priority to a specific DSCP. When the switch applies this policy to a packet, the priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. If the packet is IPv4, the packet's DSCP will be replaced by the codepoint specified in this command. (Default: For most codepoints,* **No-override**. *See figure 15-8 on page 15-52.)*

Syntax: qos < udp-port | tcp-port > < *tcp or udp port number* > dscp < *codepoint* >

Assigns a DSCP policy to packets having the specified TCP or UDP application port number and overwrites the DSCP in these packets with the assigned ***<codepoint>*** *value. This policy includes an 802.1p priority and determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. (Default:* **No-override***)*

no qos < udp-port | tcp-port > < *tcp-udp port number* >

Deletes the specified UDP or TCP port number as a QoS classifier.

show qos tcp-udp-port-priority

Displays a listing of all TCP and UDP QoS classifiers currently in the running-config file.

For example, suppose you wanted to assign these DSCP policies to the packets identified by the indicated UDP and TDP port applications:

Port Applications	DSCP Policies	
	DSCP	Priority
23-UDP	000111	7
80-TDP	000101	5
914-TDP	000010	1
1001-UDP	000010	1

1. Determine whether the DSCPs already have priority assignments, which could indicate use by existing applications. (Also, a DSCP must have a priority configured before you can assign any QoS classifiers to use it.)

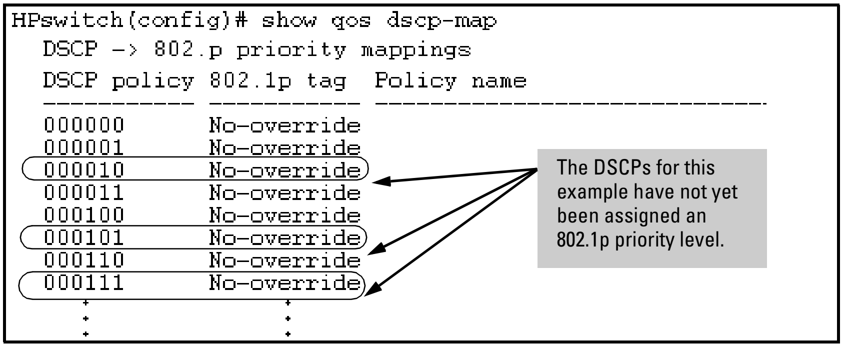

```
HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag  Policy name
  ----------- ----------- ------------------------------
  000000      No-override
  000001      No-override
  000010      No-override
  000011      No-override
  000100      No-override
  000101      No-override
  000110      No-override
  000111      No-override
      .             .
      .             .
      .             .
```

The DSCPs for this example have not yet been assigned an 802.1p priority level.

Figure 15-4. Display the Current DSCP-Map Configuration

2. Configure the DSCP policies for the codepoints you want to use.

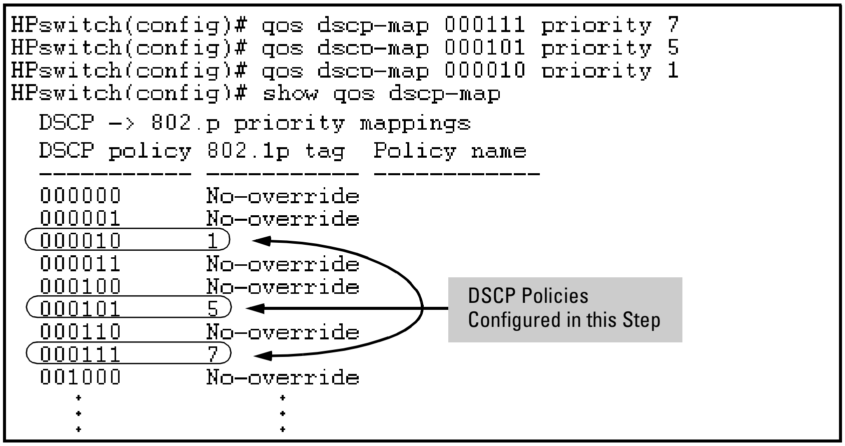

```
HPswitch(config)# qos dscp-map 000111 priority 7
HPswitch(config)# qos dscp-map 000101 priority 5
HPswitch(config)# qos dscp-map 000010 priority 1
HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag  Policy name
  ----------- ----------- ------------
  000000      No-override
  000001      No-override
  000010      1
  000011      No-override
  000100      No-override
  000101      5
  000110      No-override
  000111      7
  001000      No-override
     .             .
     .             .
     .             .
```

Figure 15-5. Assign Priorities to the Selected DSCPs

3. Assign the DSCP policies to the selected UDP/TCP port applications and display the result.

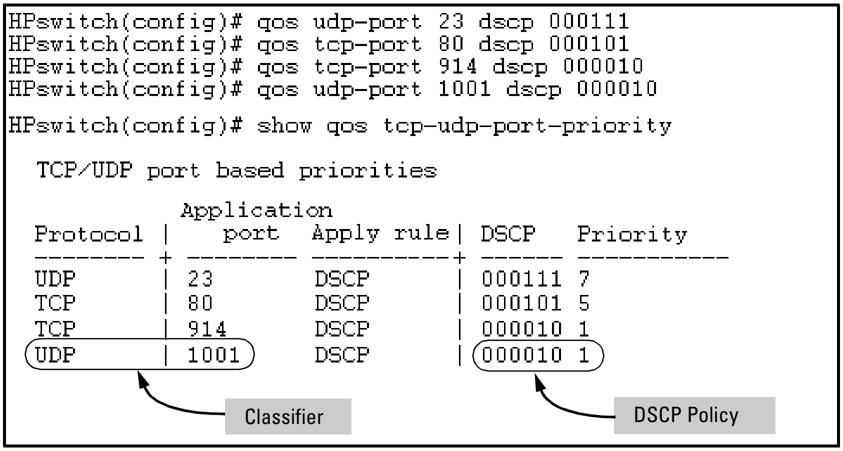

```
HPswitch(config)# qos udp-port 23 dscp 000111
HPswitch(config)# qos tcp-port 80 dscp 000101
HPswitch(config)# qos tcp-port 914 dscp 000010
HPswitch(config)# qos udp-port 1001 dscp 000010

HPswitch(config)# show qos tcp-udp-port-priority

  TCP/UDP port based priorities

            Application
  Protocol |    port   Apply rule| DSCP   Priority
  -------- + -------- ----------+ ------ -----------
  UDP      | 23       DSCP      | 000111 7
  TCP      | 80       DSCP      | 000101 5
  TCP      | 914      DSCP      | 000010 1
  UDP      | 1001     DSCP      | 000010 1
```

Figure 15-6. The Completed DSCP Policy Configuration for the Specified UDP/TCP Port Applications

The switch will now apply the DSCP policies in figure 15-6 to IP packets received in the switch with the specified UDP/TCP port applications. This means the switch will:

- Overwrite the original DSCPs in the selected packets with the new DSCPs specified in the above policies.
- Assign the 802.1p priorities in the above policies to the selected packets.

QoS IP-Device Priority

QoS Classifier Precedence: 2

The IP device option enables you to use up to 256 IP addresses (source or destination) as QoS classifiers. Where a particular device-IP address classifier has the highest precedence in the switch for traffic addressed to or from that device, then traffic received on the switch with that address is marked with the IP address classifier's configured priority level. Different IP device classifiers can have differing priority levels.

Options for Assigning Priority. Priority control options for packets carrying a specified IP address include:

- 802.1p priority
- DSCP policy (Assigning a new DSCP and an 802.1p priority; inbound packets must be IPv4.)

(For operation when other QoS classifiers apply to the same traffic, refer to "Classifiers for Prioritizing Outbound Packets" on page 15-8.)

For a given IP address, you can use only one of the above options at a time. However, for different IP addresses, you can use different options.

Assigning a Priority Based on IP Address

This option assigns an 802.1p priority to all outbound packets having the specified IP address as either a source or destination. (If both match, the priority for the IP destination address has precedence.)

Syntax: qos device-priority < *ip-address* > priority < 0 - 7 >

Configures an 802.1p priority for outbound packets having the specified IP address. This priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. (Default: Disabled)

no qos device-priority < *ip-address* >

Removes the specified IP device-priority QoS classifier and resets the priority for that VLAN to **No-override**.

show qos device-priority

Displays a listing of all IP device-priority QoS classifiers currently in the running-config file.

For example, configure and list the 802.1p priority for packets carrying the following IP addresses:

IP Address	802.1p Priority
10.28.31.1	7
10.28.31.130	5
10.28.31.100	1
10.28.31.101	1

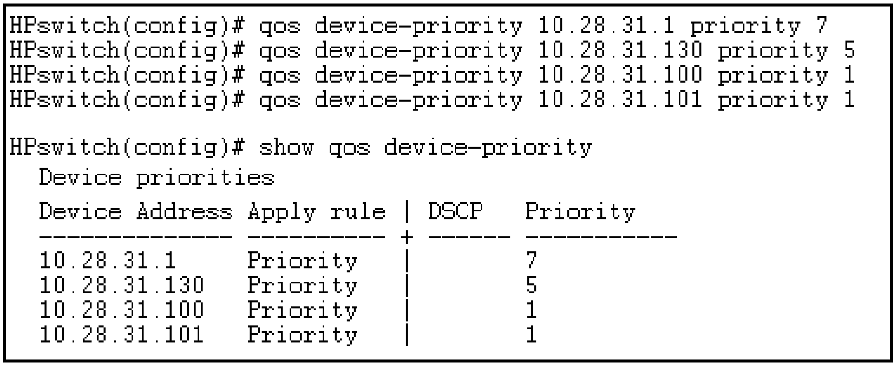

```
HPswitch(config)# qos device-priority 10.28.31.1 priority 7
HPswitch(config)# qos device-priority 10.28.31.130 priority 5
HPswitch(config)# qos device-priority 10.28.31.100 priority 1
HPswitch(config)# qos device-priority 10.28.31.101 priority 1

HPswitch(config)# show qos device-priority
  Device priorities

  Device Address Apply rule | DSCP   Priority
  -------------- ---------- + ------ -----------
  10.28.31.1     Priority   |        7
  10.28.31.130   Priority   |        5
  10.28.31.100   Priority   |        1
  10.28.31.101   Priority   |        1
```

Figure 15-7. Example of Configuring and Listing 802.1p Priority Assignments for Packets Carrying Specific IP Addresses

Assigning a DSCP Policy Based on IP Address

This option assigns a previously configured DSCP policy (codepoint and 802.1p priority) to outbound IP packets having the specified IP address (either source or destination). That is, the switch:

1. Selects an incoming IP packet on the basis of the source or destination IP address it carries.

2. Overwrites the packet's DSCP with the DSCP configured in the switch for such packets, and assigns the 802.1p priority configured in the switch for the new DSCP. (Refer to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)

3. Forwards the packet through the appropriate outbound port queue.

For more on DSCP, refer to "Terminology" on page 15-5.

Steps for Creating a Policy Based on IP Address. This procedure creates a DSCP policy for IP packets carrying the selected IP address (source or destination).

1. Identify the IP address you want to use as a classifier for assigning a DSCP policy.
2. Determine the DSCP policy for packets carrying the selected IP address:
   a. Determine the DSCP you want to assign to the selected packets. (This codepoint will be used to overwrite the DSCP carried in packets received from upstream devices.)
   b. Determine the 802.1p priority you want to assign to the DSCP.
3. Configure the DSCP policy by using **dscp-map** to configure the priority to the codepoint you selected in step 2a. (For details, refer to “Differentiated Services Codepoint (DSCP) Mapping” on page 15-51.)

Note

A codepoint must have an 802.1p priority assignment (0 - 7) before you can configure a policy for prioritizing packets by IP address. If a codepoint you want to use shows **No-override** in the **Priority** column of the DSCP map (**show qos dscp-map**), then you must assign a 0 - 7 priority before proceeding.

4. Configure the switch to assign the DSCP policy to packets with the specified IP address.

Syntax: qos dscp-map < *codepoint* > priority < 0 - 7 >

This command is optional if a priority has already been assigned to the < codepoint>. The command creates a DSCP policy by assigning an 802.1p priority to a specific DSCP. When the switch applies this policy to a packet, the priority determines the packet’s queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. If the packet is IPv4, the packet’s DSCP will be replaced by the codepoint specified in this command. (Default: For most codepoints, **No-override**. *See figure 15-8 on page 15-52.)*

Syntax: qos device-priority < *ip-address* > dscp < *codepoint* >

Assigns a DSCP policy to packets carrying the specified IP address, and overwrites the DSCP in these packets with the assigned **< *codepoint* >** *value. This policy includes an 802.1p priority and determines the packet’s queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. (Default:* **No-override***)* no qos device-priority < *ip-address* >

Deletes the specified IP address as a QoS classifier.

show qos device-priority

Displays a listing of all QoS Device Priority classifiers currently in the running-config file.

For example, suppose you wanted to assign these DSCP policies to the packets identified by the indicated IP addresses:

IP Address	DSCP Policies	
	DSCP	Priority
10.28.31.1	000111	7
10.28.31.130	000101	5
10.28.31.100	000010	1
10.28.31.101	000010	1

1. Determine whether the DSCPs already have priority assignments, which could indicate use by existing applications. This is not a problem as long as the configured priorities are acceptable for all applications using the same DSCP. (Refer to the “Note On Changing a Priority Setting” on page 15-54. Also, a DSCP must have a priority configured before you can assign any QoS classifiers to use it.)

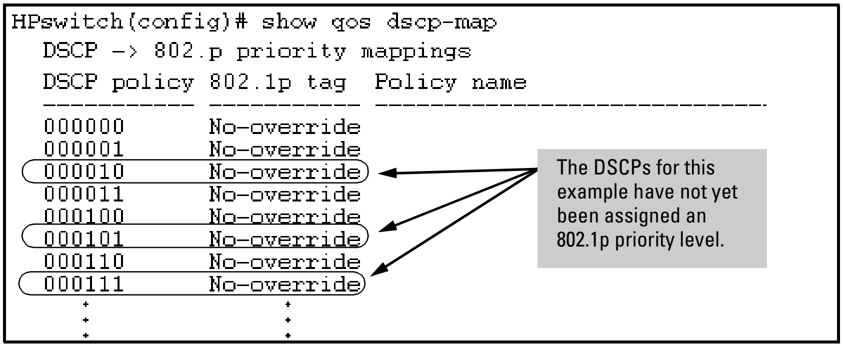

```
HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings

  DSCP policy 802.1p tag  Policy name
  ----------- ----------- -----------------------------
  000000      No-override
  000001      No-override
  000010      No-override
  000011      No-override
  000100      No-override
  000101      No-override
  000110      No-override
  000111      No-override
      .            .
      .            .
      .            .
```

The DSCPs for this example have not yet been assigned an 802.1p priority level.

Figure 15-8. Display the Current DSCP-Map Configuration

2. Configure the priorities for the DSCPs you want to use.

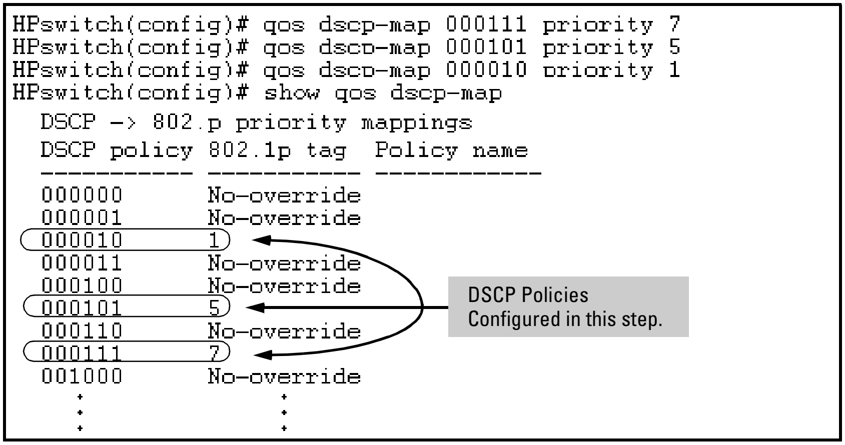

```
HPswitch(config)# qos dscp-map 000111 priority 7
HPswitch(config)# qos dscp-map 000101 priority 5
HPswitch(config)# qos dscp-map 000010 priority 1
HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag  Policy name
  ----------- ----------- ------------
  000000      No-override
  000001      No-override
  000010      1
  000011      No-override
  000100      No-override
  000101      5
  000110      No-override
  000111      7
  001000      No-override
     .             .
     .             .
     .             .
```

DSCP Policies Configured in this step.

Figure 15-9. Assigning 802.1p Priorities to the Selected DSCPs

3. Assign the DSCP policies to the selected device IP addresses and display the result.

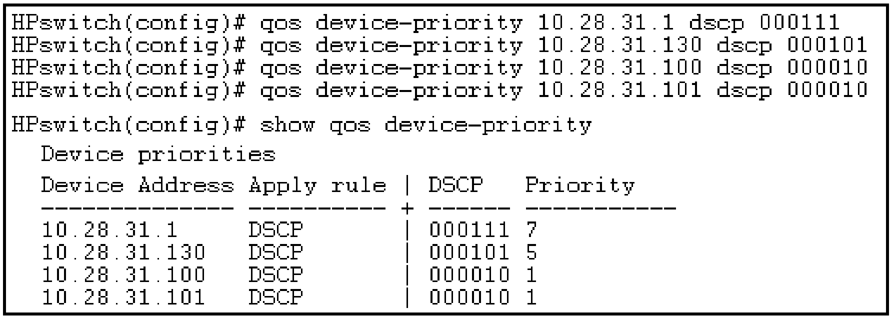

```
HPswitch(config)# qos device-priority 10.28.31.1 dscp 000111
HPswitch(config)# qos device-priority 10.28.31.130 dscp 000101
HPswitch(config)# qos device-priority 10.28.31.100 dscp 000010
HPswitch(config)# qos device-priority 10.28.31.101 dscp 000010

HPswitch(config)# show qos device-priority
  Device priorities
  Device Address Apply rule | DSCP   Priority
  -------------- ---------- + ------ -----------
  10.28.31.1     DSCP       | 000111 7
  10.28.31.130   DSCP       | 000101 5
  10.28.31.100   DSCP       | 000010 1
  10.28.31.101   DSCP       | 000010 1
```

Figure 15-10. The Completed Device-Priority/Codepoint Configuration

The switch will now apply the DSCP policies in figure 15-9 to packets received on the switch with the specified IP addresses (source or destination). This means the switch will:

- Overwrite the original DSCPs in the selected packets with the new DSCPs specified in the above policies.
- Assign the 802.1p priorities in the above policies to the appropriate packets.

QoS IP Type-of-Service (ToS) Policy and Priority

QoS Classifier Precedence: 3

This feature applies only to IP traffic and performs either of the following:

- **ToS IP-Precedence Mode:** All IP packets generated by upstream devices and applications include precedence bits in the ToS byte. Using this mode, the switch uses these bits to compute and assign the corresponding 802.1p priority.
- **ToS Differentiated Services (Diffserv) Mode:** This mode requires knowledge of the codepoints set in IP packets by the upstream devices and applications. It uses the ToS codepoint in IP packets coming from upstream devices and applications to assign 802.1p priorities to the packets. You can use this option to do both of the following:
  - **Assign a New Prioritization Policy:** A "policy" includes both a codepoint and a corresponding 802.1p priority. This option selects an incoming IP packet on the basis of its codepoint and assigns a new codepoint and corresponding 802.1p priority. (Use the **qos dscp-map** command to specify a priority for any codepoint—page 15-51.)
  - **Assign an 802.1p Priority:** This option reads the DSCP of an incoming IP packet and, without changing this codepoint, assigns the 802.1p priority to the packet, as configured in the DSCP Policy Table (page 15-51). This means that a priority value of 0 - 7 must be configured for a DSCP before the switch will attempt to perform a QoS match on the packet's DSCP bits.

  Before configuring the ToS Diffserv mode, you must use the **dscp-map** command to configure the desired 802.1p priorities for the codepoints you want to use for either option. This command is illustrated in the following examples and is described under "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.

Unless IP-Precedence mode and Diffserv mode are both disabled (the default setting), enabling one automatically disables the other. *For more on ToS operation, refer to "Details of QoS IP Type-of-Service" on page 15-34.*

Assigning an 802.1p Priority to IP Packets on the Basis of the ToS Precedence Bits

If a device or application upstream of the switch sets the precedence bits in the ToS byte of IP packets, you can use this feature to apply that setting for prioritizing packets for outbound port queues. If the outbound packets are in a tagged VLAN, this priority is carried as an 802.1p value to the adjacent downstream devices.

IP Precedence Syntax: qos type-of-service ip-precedence

Causes the switch to automatically assign an 802.1p priority to all IP packets by computing each packet's 802.1p priority from the precedence bits the packet carries. This priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. (ToS IP Precedence Default: Disabled)

no qos type-of-service

Disables all ToS classifier operation, including prioritization using the precedence bits.

show qos type-of-service

When ip-precedence is enabled (or if neither ToS option is configured), shows the ToS configuration status. If diff-services is enabled, lists codepoint data as described under "Assigning a DSCP Policy on the Basis of the DSCP in IP Packets Received from Upstream Devices" on page 15-31.

With this option, prioritization of outbound packets relies on the IP-Precedence bit setting that IP packets carry with them from upstream devices and applications. To configure and verify this option:

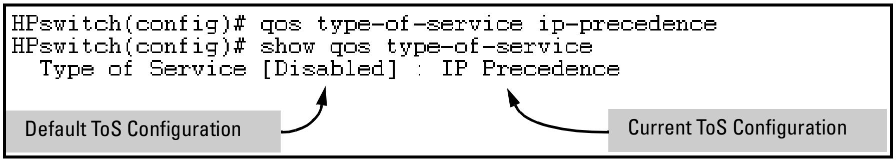

```
HPswitch(config)# qos type-of-service ip-precedence
HPswitch(config)# show qos type-of-service
  Type of Service [Disabled] : IP Precedence
```

Figure 15-11. Example of Enabling ToS IP-Precedence Prioritization

To replace this option with the ToS diff-services option, just configure **diff-services** as described below, which automatically disables IP-Precedence. To disable IP-Precedence without enabling the diff-services option, use this command:

```
HPswitch(config)# no qos type-of-service k
```

Assigning an 802.1p Priority to IP Packets on the Basis of Incoming DSCP

One of the best uses for this option is on an interior switch where you want to honor (continue) a policy set on an edge switch. That is, it enables you to select incoming packets having a specific DSCP and forward these packets with the desired 802.1p priority. For example, if an edge switch "A" marks all packets received on port A5 with a particular DSCP, you can configure a downstream (interior) switch "B" to handle such packets with the desired priority (regardless of whether 802.1Q tagged VLANs are in use).

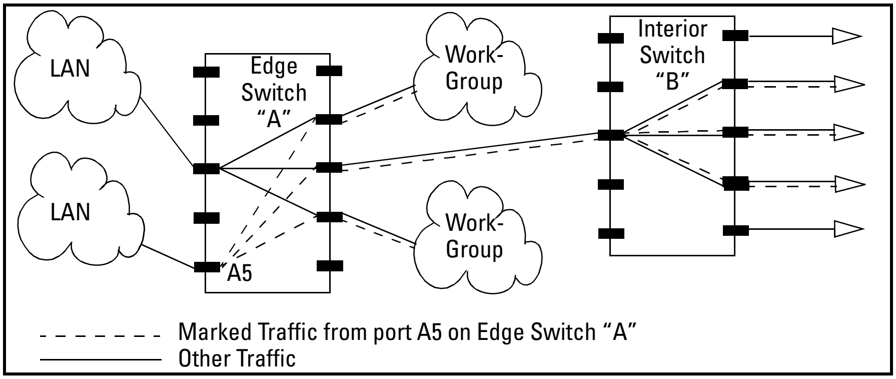

Figure 15-12. Interior Switch "B" Honors the Policy Established in Edge Switch "A"

To do so, assign the desired 802.1p priority to the same codepoint that the upstream or edge switch assigns to the selected packets. When the downstream switch receives an IP packet carrying one of these codepoints, it assigns the configured priority to the packet and sends it out the appropriate priority queue. (The packet retains the codepoint it received from the upstream or edge switch). You can use this option concurrently with the diffserv DSCP Policy option (described later in this section), as long as the DSCPs specified in the two options do not match.

Operating Notes

Different applications may use the same DSCP in their IP packets. Also, the same application may use multiple DSCPs it the application originates on different clients, servers, or other devices. Using an edge switch enables you to select the packets you want and mark them with predictable DSCPs that can be used by downstream switches to honor policies set in the edge switch.

When enabled, the switch applies direct 802.1p prioritization to all packets having codepoints that meet these criteria:

- The codepoint is configured with an 802.1p priority in the DSCP table. (Codepoints configured with **No-override** are not used.)
- The codepoint is not configured for a new DSCP policy assignment.

Thus, the switch does not allow the same incoming codepoint (DSCP) to be used simultaneously for directly assigning an 802.1p priority and also assigning a DSCP policy. For a given incoming codepoint, if you configure one option and then the other, the second overwrites the first.

To use this option:

1. Identify a DSCP used to set a policy in packets received from an upstream or edge switch.
2. Determine the 802.1p priority (0 - 7) you want to apply to packets carrying the identified DSCP. (You can either maintain the priority assigned in the upstream or edge switch, or assign a new priority.)
3. Use **qos dscp-map <** ***codepoint*** **> priority < 0 - 7 >** to assign the 802.1p priority you want to the specified DSCP. (For more on this topic, refer to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)
4. Enable **diff-services**

.

Syntax: qos type-of-service diff-services

Causes the switch to read the DSCP of an incoming IP packet and, when a match occurs, assign a corresponding 802.1p priority, as configured in the switch's DSCP table (page 15-52).

no qos type-of-service

Disables all ToS classifier operation.

no qos dscp-map < *codepoint* >

Disables direct 802.1p priority assignment to packets carrying the < codepoint> *by reconfiguring the codepoint priority assignment in the DSCP table to* **No-override**. *Note that if this codepoint is in use as a DSCP policy for another diffserv codepoint, you must disable or redirect the other diffserv codepoint's DSCP policy before you can disable or change the codepoint. For example, in figure 15-13 you cannot change the priority for the 000000 codepoint until you redirect the DSCP policy for 000001 away from using 000000 as a policy. (Refer to "Note On Changing a Priority Setting" on page 15-54. Refer also to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)*

show qos type-of-service

Displays current Type-of-Service configuration. In diffserv mode it also shows the current direct 802.1p assignments and the current DSCP assignments covered later in this section.

For example, an edge switch "A" in an untagged VLAN assigns a DSCP of 000110 on IP packets it receives on port A6, and handles the packets with high priority (7). When these packets reach interior switch "B" you want the switch to handle them with the same high priority. To enable this operation you would configure an 802.1p priority of 7 for packets received with a DSCP of **000110**, and then enable **diff-services**:

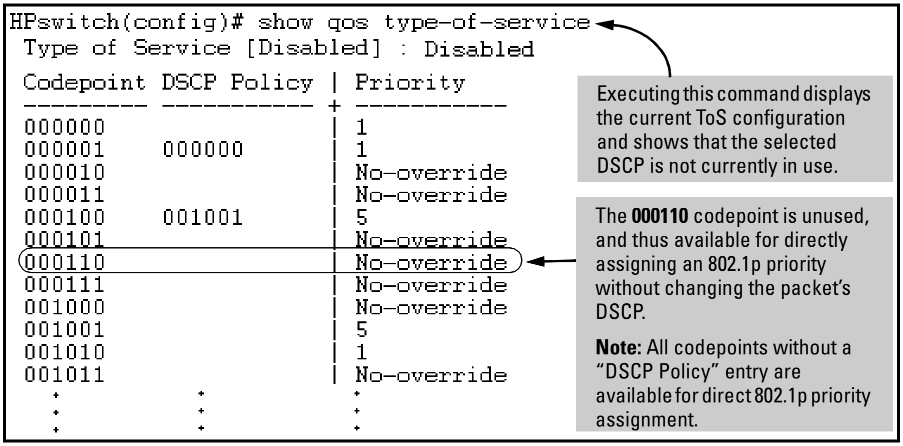

```
HPswitch(config)# show qos type-of-service
 Type of Service [Disabled] : Disabled

 Codepoint DSCP Policy | Priority
 --------- ----------- + -----------
 000000                | 1
 000001    000000      | 1
 000010                | No-override
 000011                | No-override
 000100    001001      | 5
 000101                | No-override
 000110                | No-override
 000111                | No-override
 001000                | No-override
 001001                | 5
 001010                | 1
 001011                | No-override
   .          .          .
   .          .          .
   .          .          .
```

Executing this command displays the current ToS configuration and shows that the selected DSCP is not currently in use.

The **000110** codepoint is unused, and thus available for directly assigning an 802.1p priority without changing the packet's DSCP.

Note: All codepoints without a "DSCP Policy" entry are available for direct 802.1p priority assignment.

Figure 15-13. Example Showing Codepoints Available for Direct 802.1p Priority Assignments

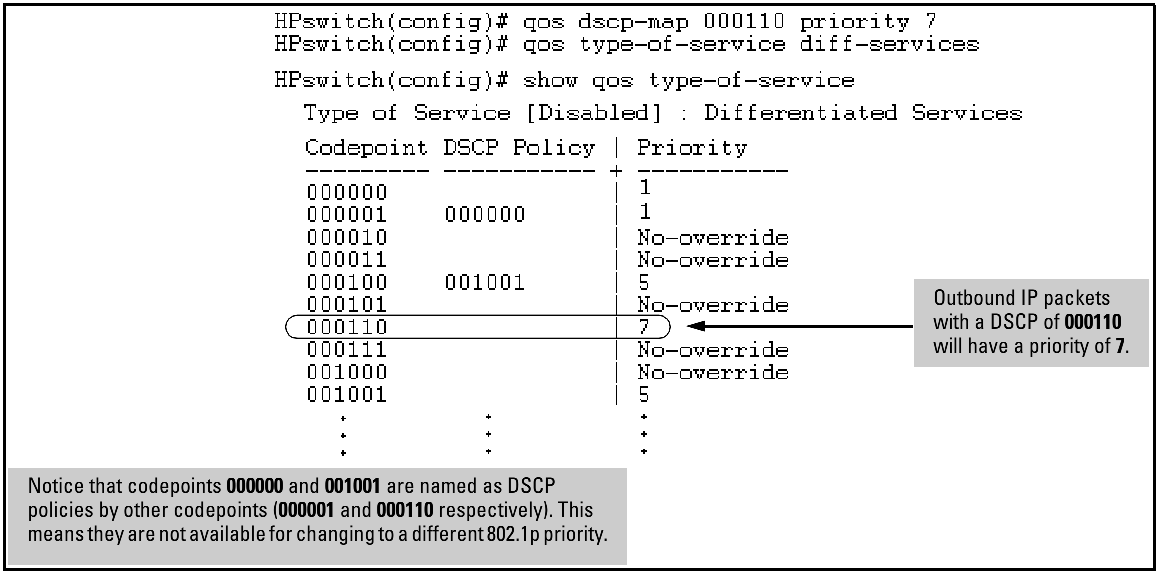

```
HPswitch(config)# qos dscp-map 000110 priority 7
HPswitch(config)# qos type-of-service diff-services

HPswitch(config)# show qos type-of-service
  Type of Service [Disabled] : Differentiated Services

  Codepoint DSCP Policy | Priority
  --------- ----------- + -----------
  000000                | 1
  000001    000000      | 1
  000010                | No-override
  000011                | No-override
  000100    001001      | 5
  000101                | No-override
  000110                | 7
  000111                | No-override
  001000                | No-override
  001001                | 5
    .          .          .
    .          .          .
    .          .          .
```

Outbound IP packets with a DSCP of **000110** will have a priority of **7**.

Notice that codepoints **000000** and **001001** are named as DSCP policies by other codepoints (**000001** and **000110** respectively). This means they are not available for changing to a different 802.1p priority.

Figure 15-14. Example of a Type-of-Service Configuration Enabling Both Direct 802.1p Priority Assignment and DSCP Policy Assignment

Assigning a DSCP Policy on the Basis of the DSCP in IP Packets Received from Upstream Devices

The preceding section describes how to forward a policy set by an edge (or upstream) switch. This option changes a DSCP policy in an outbound IP packet by changing its IP ToS codepoint and applying the priority associated with the new codepoint. (A DSCP policy consists of a differentiated services codepoint and an associated 802.1p priority.) You can use this option concurrently with the diffserv 802.1p priority option (above), as long as the DSCPs specified in the two options do not match.

To use this option to configure a change in policy:

1. Identify a DSCP used to set a policy in packets received from an upstream or edge switch.
2. Create a new policy by using **qos dscp-map < *codepoint* > priority < 0 - 7 >** to configure an 802.1p priority for the codepoint you will use to overwrite the DSCP the packet carries from upstream. (For more on this topic, refer to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)
3. Use **qos type-of-service diff-services < *incoming-DSCP* > dscp < *outgoing-DSCP* >** to change the policy on packets coming from the edge or upstream switch with the specified incoming DSCP.

   (Figure 15-12 on page 15-27 illustrates this scenario.)

Syntax: qos type-of-service diff-services

Enables ToS diff-services.

qos type-of-service diff-services < *current-codepoint* > dscp < *new-codepoint* >

Configures the switch to select an incoming IP packet carrying the <current-codepoint> and then use the <new-codepoint> to assign a new, previously configured DSCP policy to the packet. The policy overwrites the <current-codepoint> with the < new-codepoint > and assigns the 802.1p priority specified by the policy. (Use the **qos dscp-map** *command to define the priority for the DSCPs—page 15-51.)*

no qos type-of-service

Disables all ToS classifier operation. Current ToS DSCP policies and priorities remain in the configuration and will become available if you re-enable ToS diff-services.

no qos type-of-service [diff-services < *codepoint* >]

Deletes the DSCP policy assigned to the < codepoint > and returns the < codepoint > to the 802.1p priority setting it had before the DSCP policy was assigned. (This will be either a value from 0 - 7 or **No-override***.)*

show qos type-of-service

Displays a listing of codepoints, with any corresponding DSCP policy re-assignments for outbound packets. Also lists the (802.1p) priority for each codepoint that does not have a DSCP policy assigned to it.

For example, suppose you want to configure the following two DSCP policies for packets received with the indicated DSCPs.

Received DSCP	Policy DSCP	802.1p Priority	Policy Name (Optional)
001100	000010	6	Level 6
001101	000101	4	Level 4

1. Determine whether the DSCPs already have priority assignments, which could indicate use by existing applications. This is not a problem as long as the configured priorities are acceptable for all applications using the same DSCP. (Refer to the “Note On Changing a Priority Setting” on page 15-54. Also, a DSCP must have a priority configured before you can assign any QoS classifiers to use it.)

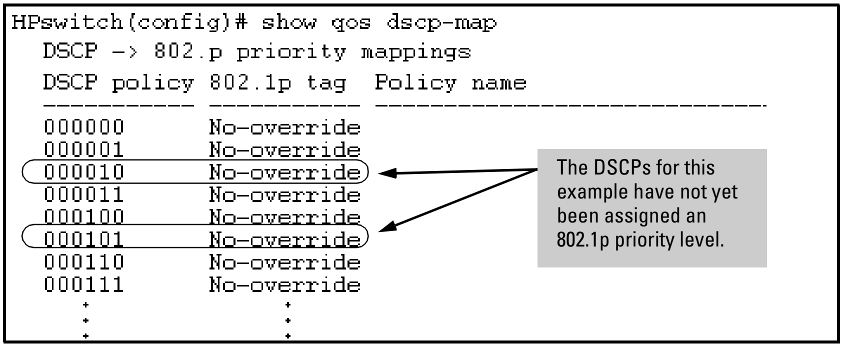

```
HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag  Policy name
  ----------- ----------- -----------------------------
  000000      No-override
  000001      No-override
  000010      No-override
  000011      No-override
  000100      No-override
  000101      No-override
  000110      No-override
  000111      No-override
     .              .
     .              .
     .              .
```

The DSCPs for this example have not yet been assigned an 802.1p priority level.

Figure 15-15. Display the Current DSCP-Map Configuration

2. Configure the policies in the DSCP table:

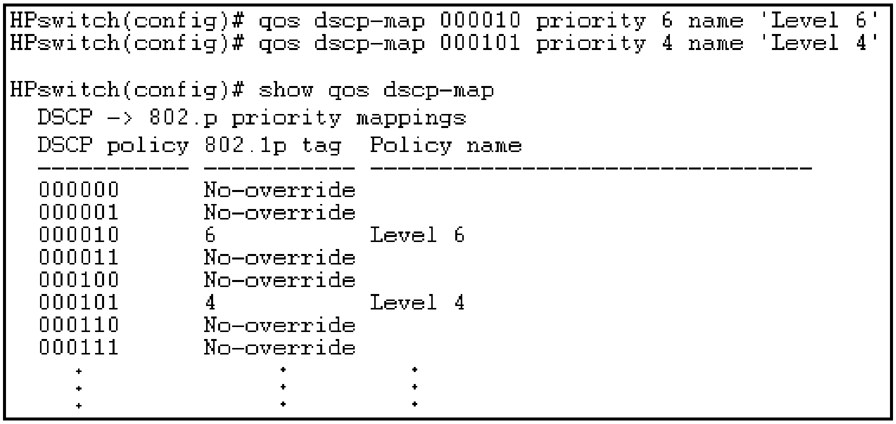

```
HPswitch(config)# qos dscp-map 000010 priority 6 name 'Level 6'
HPswitch(config)# qos dscp-map 000101 priority 4 name 'Level 4'

HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag  Policy name
  ----------- ----------- -------------------------------
  000000      No-override
  000001      No-override
  000010      6           Level 6
  000011      No-override
  000100      No-override
  000101      4           Level 4
  000110      No-override
  000111      No-override
     .              .         .
     .              .         .
     .              .         .
```

Figure 15-16. Example of Policies Configured (with Optional Names) in the DSCP Table

3. Assign the policies to the codepoints in the selected packet types.

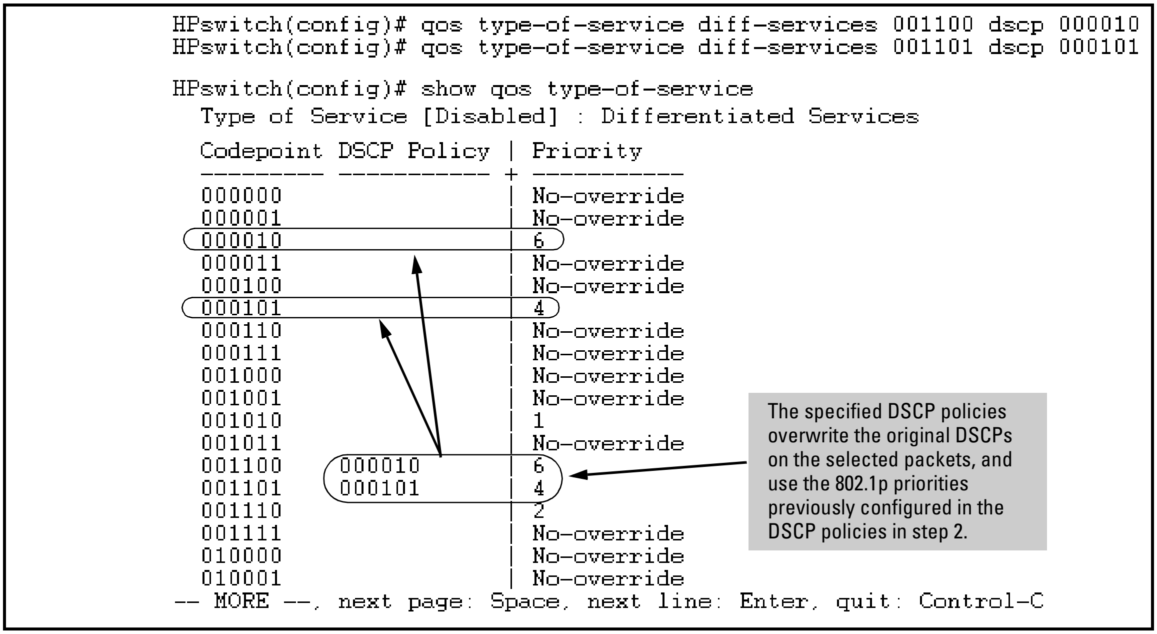

```
HPswitch(config)# qos type-of-service diff-services 001100 dscp 000010
HPswitch(config)# qos type-of-service diff-services 001101 dscp 000101

HPswitch(config)# show qos type-of-service
  Type of Service [Disabled] : Differentiated Services

  Codepoint DSCP Policy | Priority
  --------- ----------- + -----------
  000000                | No-override
  000001                | No-override
  000010                | 6
  000011                | No-override
  000100                | No-override
  000101                | 4
  000110                | No-override
  000111                | No-override
  001000                | No-override
  001001                | No-override
  001010                | 1
  001011                | No-override
  001100    000010      | 6
  001101    000101      | 4
  001110                | 2
  001111                | No-override
  010000                | No-override
  010001                | No-override
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

The specified DSCP policies overwrite the original DSCPs on the selected packets, and use the 802.1p priorities previously configured in the DSCP policies in step 2.

Figure 15-17. Example of Policy Assignment to Outbound Packets on the Basis of the DSCP in the Packets Received from Upstream Devices

Details of QoS IP Type-of-Service

IP packets include a Type of Service (ToS) byte. The ToS byte includes:

- **A Differentiated Services Codepoint (DSCP):** This element is comprised of the upper six bits of the ToS byte). There are 64 possible codepoints. In the switch's default **qos** configuration, some codepoints have default 802.1p priority settings for Assured-Forwarding and Expedited Forwarding, while others are unused (and listed with **No-override** for a Priority). Using the **qos dscp map** command, you can configure the switch to assign different prioritization policies to IP packets having different codepoints. As an alternative, you can configure the switch to assign a new codepoint to an IP packet, along with a corresponding 802.1p priority (0-7). To use this option in the simplest case, you would:

  a. Configure a specific DSCP with a specific priority in an edge switch.

  b. Configure the switch to mark a specific type of inbound traffic with that DSCP (and thus create a policy for that traffic type).

  c. Configure the internal switches in your LAN to honor the policy.

  (For example, you could configure an edge switch to assign a codepoint of 000001 to all packets received from a specific VLAN, and then handle all traffic with that codepoint at high priority.)

  For a codepoint listing and the commands for displaying and changing the DSCP Policy table, refer to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.

- **Precedence Bits:** This element is a subset of the DSCP and is comprised of the upper three bits of the ToS byte. When configured to do so, the switch uses the precedence bits to determine a priority for handling the associated packet. (The switch does not change the setting of the precedence bits.) Using the ToS Precedence bits to prioritize IP packets relies on priorities set in upstream devices and applications.

Figure 15-18 shows an example of the ToS byte in the header for an IP packet, and illustrates the diffserv bits and precedence bits in the ToS byte. (Note that the Precedence bits are a subset of the Differentiated Services bits.)

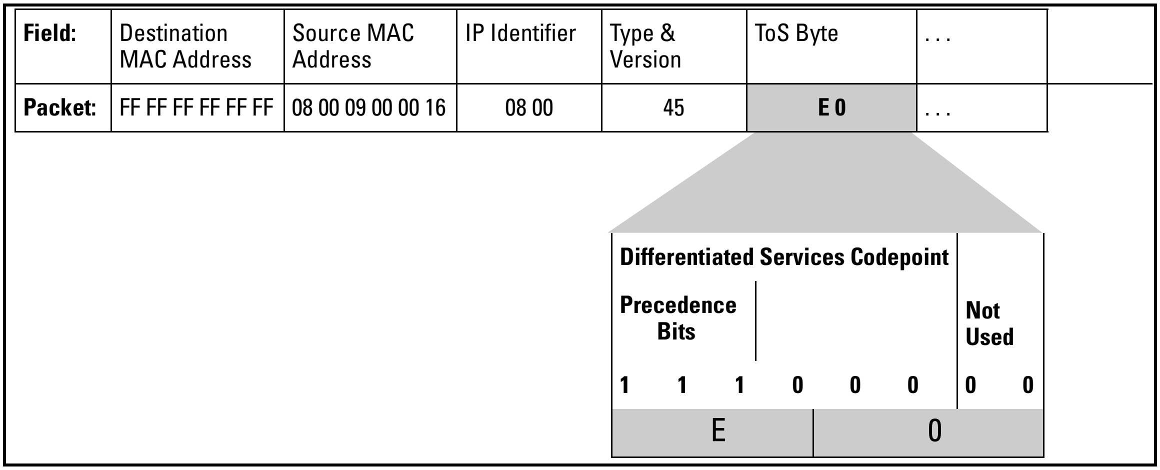

Field:	Destination MAC Address	Source MAC Address	IP Identifier	Type & Version	ToS Byte	. . .
Packet:	FF FF FF FF FF FF	08 00 09 00 00 16	08 00	45	**E 0**	. . .

Figure 15-18. The ToS Codepoint and Precedence Bits

Table 16-7.How the Switch Uses the ToS Configuration

Outbound Port	**ToS Option:**	
	IP Precedence (Value = 0 - 7)	**Differentiated Services**
IP Packet in an Untagged VLAN	Depending on the value of the IP Precedence bits in the packet's ToS field, the packet will go to one of four outbound port queues in the switch: 1 - 2 = low priority 0 - 3 = normal priority 4 - 5 = high priority 6 - 7 = high priority	For a given packet carrying a ToS codepoint that the switch has been configured to detect: • Change the codepoint according to the configured policy and assign the 802.1p priority specified for the new codepoint in the DSCP Policy Table (page 15-51). • Do not change the codepoint, but assign the 802.1p priority specified for the existing codepoint in the DSCP Policy Table (page 15-51). Depending on the 802.1p priority used, the packet will leave the switch through one of the following queues: 1 - 2 = low priority 0 - 3 = normal priority 4 - 5 = high priority 6 - 7 = high priority If **No-override** (the default) has been configured for a specified codepoint, then the packet is not prioritized by ToS and, by default, is sent to the "normal priority" queue.
IP Packet in a Tagged VLAN	Same as above, plus the IP Precedence value (0 - 7) will be used to set a corresponding 802.1p priority in the VLAN tag carried by the packet to the next downstream device.	Same as above, plus the Priority value (0 - 7) will be used to set a corresponding 802.1p priority in the VLAN tag carried by the packet to the next downstream device. Where **No-override** is the assigned priority, the VLAN tag carries a "0" (normal priority) 802.1p setting.

QoS LAN Protocol Priority

QoS Classifier Precedence: 4

The QoS protocol option enables you to use these protocols as QoS classifiers:

- IP
- IPX
- ARP
- DEC_LAT
- Appletalk
- SNA
- Netbeui

Options for Assigning Priority. Priority control for the LAN protocol classifier includes assigning only the 802.1p priority. The switch does not use this classifier for assigning DSCP-based priority.

Assigning a Priority Based on LAN Protocol

When QoS on the switch is configured with a LAN protocol as the highest-precedence classifier and the switch receives traffic carrying that protocol, then this traffic is assigned the priority configured for this classifier. (For operation when other QoS classifiers apply to the same traffic, refer to "Classifiers for Prioritizing Outbound Packets" on page 15-8.)

Syntax: qos protocol

< ip | ipx | arp | dec_lat | appletalk | sna | netbeui > priority < 0 - 7 >

Configures an 802.1p priority for outbound packets having the specified protocol. This priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. You can configure one QoS classifier for each protocol type. (Default: **No-override***)*

no qos protocol

< ip | ipx | arp | dec_lat | appletalk | sna | netbeui >

Disables use of the specified protocol as a QoS classifier and resets the protocol priority to **No-override**.

show qos protocol

Lists the QoS protocol classifiers with their priority settings.

For example:

1. Configure QoS protocol classifiers with IP at 0 (normal), ARP at 5 (medium), and AppleTalk at 7 (high) and display the QoS protocol configuration.
2. Disable the QoS IP protocol classifier, downgrade the ARP priority to 4, and again display the QoS protocol configuration.

Figure 15-19 shows the command sequence and displays for the above steps.

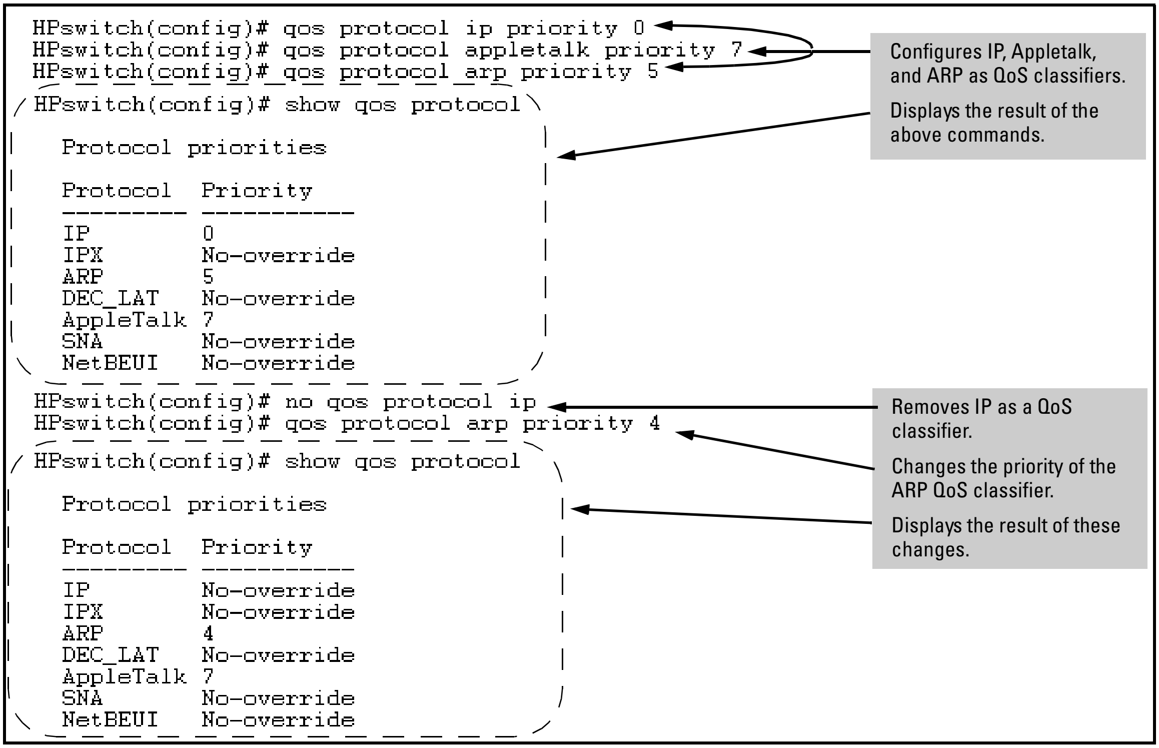

```
HPswitch(config)# qos protocol ip priority 0
HPswitch(config)# qos protocol appletalk priority 7
HPswitch(config)# qos protocol arp priority 5
HPswitch(config)# show qos protocol

  Protocol priorities

  Protocol  Priority
  --------- -----------
  IP        0
  IPX       No-override
  ARP       5
  DEC_LAT   No-override
  AppleTalk 7
  SNA       No-override
  NetBEUI   No-override
HPswitch(config)# no qos protocol ip
HPswitch(config)# qos protocol arp priority 4
HPswitch(config)# show qos protocol

  Protocol priorities

  Protocol  Priority
  --------- -----------
  IP        No-override
  IPX       No-override
  ARP       4
  DEC_LAT   No-override
  AppleTalk 7
  SNA       No-override
  NetBEUI   No-override
```

Configures IP, Appletalk, and ARP as QoS classifiers.

Displays the result of the above commands.

Removes IP as a QoS classifier.

Changes the priority of the ARP QoS classifier.

Displays the result of these changes.

Figure 15-19. Adding, Displaying, Removing, and Changing QoS Protocol Classifiers

QoS VLAN-ID (VID) Priority

QoS Classifier Precedence: 5

The QoS protocol option enables you to use up to 256 VLAN-IDs as QoS classifiers. Where a particular VLAN-ID classifier has the highest precedence in the switch for traffic in that VLAN, then traffic received in that VLAN is marked with the VID classifier's configured priority level. Different VLAN-ID classifiers can have differing priority levels.

Options for Assigning Priority. Priority control options for packets carrying a specified VLAN-ID include:

- 802.1p priority
- DSCP policy (Assigning a new DSCP and an associated 802.1p priority; inbound packets must be IPv4.)

(For operation when other QoS classifiers apply to the same traffic, refer to "Classifiers for Prioritizing Outbound Packets" on page 15-8.)

Note

QoS with VID priority applies to static VLANs only, and applying QoS to dynamic VLANs created by GVRP operation is not supported. A VLAN must exist while a subject of a QoS configuration, and eliminating a VID from the switch causes the switch to clear any QoS features configured for that VID.

Assigning a Priority Based on VLAN-ID Only

This option assigns a priority to all outbound packets having the specified VLAN-ID (VID). You can configure this option by either specifying the VLAN-ID ahead of the **qos** command or moving to the VLAN context for the VLAN you want to configure for priority.

Syntax: vlan < *vlan-id* > qos priority < 0 - 7 >

Configures an 802.1p priority for outbound packets belonging to the specified VLAN. This priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. You can configure one QoS classifier for each VLAN-ID. (Default: **No-override***)* no vlan < *vlan-id* > qos

Removes the specified VLAN-ID as a QoS classifier and resets the priority for that VLAN to **No- override**.

show qos vlan-priority

Displays a listing of the QoS VLAN-ID classifiers currently in the running-config file, with their priority data.

1. For example, suppose that you have the following VLANs configured on the switch and want to prioritize them as shown:

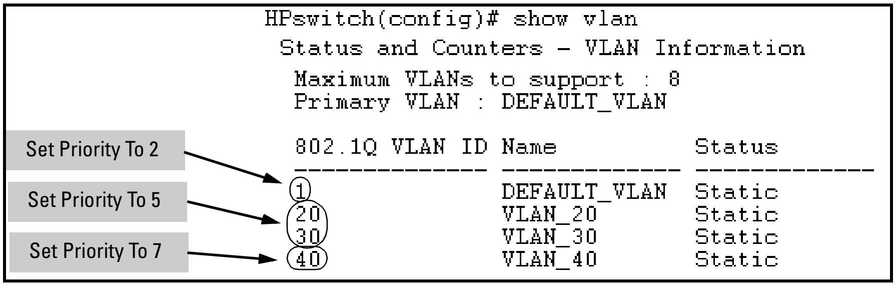

```
HPswitch(config)# show vlan
 Status and Counters - VLAN Information
  Maximum VLANs to support : 8
  Primary VLAN : DEFAULT_VLAN

  802.1Q VLAN ID Name           Status
  -------------- -------------- -------------
  1              DEFAULT_VLAN   Static
  20             VLAN_20        Static
  30             VLAN_30        Static
  40             VLAN_40        Static
```

Figure 15-20. Example of a List of VLANs Available for QoS Prioritization

2. You would then execute the following commands to prioritize the VLANs by VID:

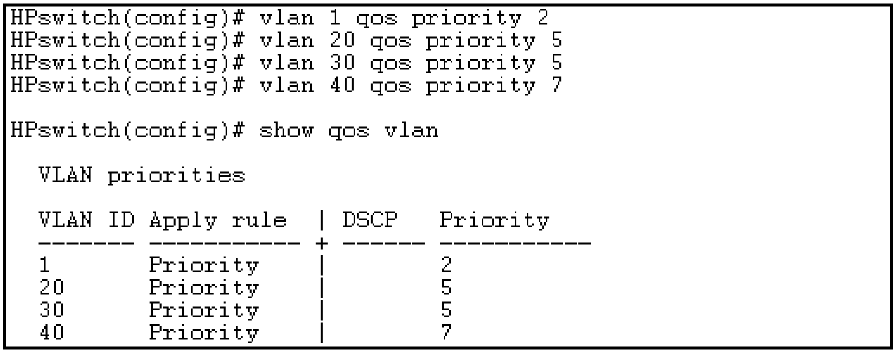

```
HPswitch(config)# vlan 1 qos priority 2
HPswitch(config)# vlan 20 qos priority 5
HPswitch(config)# vlan 30 qos priority 5
HPswitch(config)# vlan 40 qos priority 7

HPswitch(config)# show qos vlan

 VLAN priorities

 VLAN ID Apply rule  | DSCP   Priority
 ------- ----------- + ------ -----------
 1       Priority    |        2
 20      Priority    |        5
 30      Priority    |        5
 40      Priority    |        7
```

Figure 15-21. Configuring and Displaying QoS Priorities on VLANs

If you then decided to remove VLAN_20 from QoS prioritization:

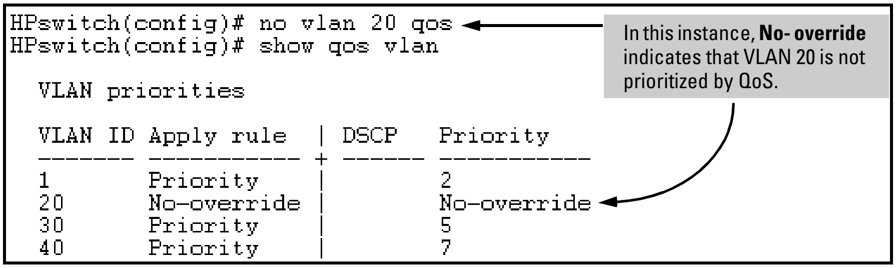

```
HPswitch(config)# no vlan 20 qos
HPswitch(config)# show qos vlan

  VLAN priorities

  VLAN ID Apply rule   | DSCP   Priority
  ------- ----------- + ------ -----------
  1       Priority    |        2
  20      No-override |        No-override
  30      Priority    |        5
  40      Priority    |        7
```

In this instance, **No- override** indicates that VLAN 20 is not prioritized by QoS.

Figure 15-22. Returning a QoS-Prioritized VLAN to "No-override" Status

Assigning a DSCP Policy Based on VLAN-ID (VID)

This option assigns a previously configured DSCP policy (codepoint and 802.1p priority) to outbound IP packets having the specified VLAN-ID (VID). That is, the switch:

1. Selects an incoming IP packet on the basis of the VLAN-ID it carries.
2. Overwrites the packet's DSCP with the DSCP configured in the switch for such packets.
3. Assigns the 802.1p priority configured in the switch for the new DSCP. (Refer to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)
4. Forwards the packet through the appropriate outbound port queue.

For more on DSCP, refer to "Terminology" on page 15-5.

Steps for Creating a Policy Based on VLAN-ID Classifier.

1. Determine the VLAN-ID classifier to which you want to assign a DSCP policy.
2. Determine the DSCP policy for packets carrying the selected VLAN-ID:
   a. Determine the DSCP you want to assign to the selected packets. (This codepoint will be used to overwrite the DSCP carried in packets received from upstream devices.)
   b. Determine the 802.1p priority you want to assign to the DSCP.
3. Configure the DSCP policy by using **qos dscp-map** to configure the priority for each codepoint. (For details, see the example later in this section, and to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)

Note

A codepoint must have an 802.1p priority (0 - 7) before you can configure the codepoint for use in prioritizing packets by VLAN-ID. If a codepoint you want to use shows **No-override** in the **Priority** column of the DSCP Policy table (**show qos dscp-map**), then assign a priority before proceeding.

4. Configure the switch to assign the DSCP policy to packets with the specified VLAN-ID.

Syntax: qos dscp-map < *codepoint* > priority < 0 - 7 >

This command is optional if a priority has already been assigned to the < codepoint>*. The command creates a DSCP policy by assigning an 802.1p priority to a specific DSCP. When the switch applies this priority to a packet, the priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. If the packet is IPv4, the packet's DSCP will be replaced by the codepoint specified in this command. (Default: For most codepoints,* **No-override***. See figure 15-8 on page 15-52.)*

Syntax: vlan < *vid* > qos dscp < *codepoint* >

Assigns a DSCP policy to packets carrying the specified IP address, and overwrites the DSCP in these packets with the assigned < **codepoint** *> value. This policy includes an 802.1p priority and determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. (Default:* **No-override***)*

no vlan < *vid* > qos

Removes QoS classifier for the specified VLAN.

show qos device-priority

Displays a listing of all QoS VLAN-ID classifiers currently in the running-config file.

For example, suppose you wanted to assign this set of priorities:

VLAN-ID	DSCP	Priority
40	000111	7
30	000101	5
20	000010	1
1	000010	1

1. Determine whether the DSCPs already have priority assignments, which could indicate use by existing applications. This is not a problem as long as the configured priorities are acceptable for all applications using the same DSCP. (Refer to the "Note On Changing a Priority Setting" on page 15-54. Also, a DSCP must have a priority configured before you can assign any QoS classifiers to use it.)

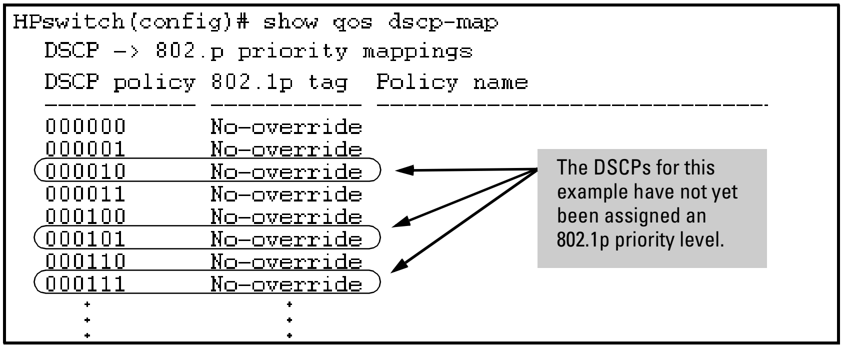

```
HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag  Policy name
  ----------- ----------- -------------------------------
  000000      No-override
  000001      No-override
  000010      No-override
  000011      No-override
  000100      No-override
  000101      No-override
  000110      No-override
  000111      No-override
     :             :
```

The DSCPs for this example have not yet been assigned an 802.1p priority level.

Figure 15-23. Display the Current Configuration in the DSCP Policy Table

2. Configure the priorities for the DSCPs you want to use.

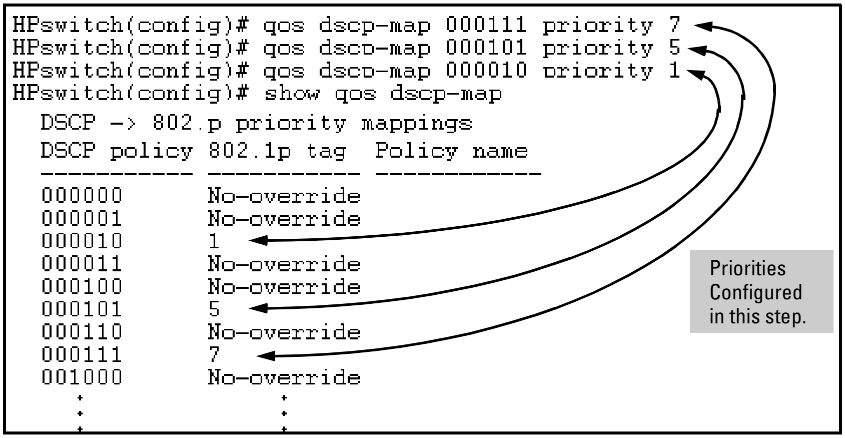

```
HPswitch(config)# qos dscp-map 000111 priority 7
HPswitch(config)# qos dscp-map 000101 priority 5
HPswitch(config)# qos dscp-map 000010 priority 1
HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag  Policy name
  ----------- ----------- ------------
  000000      No-override
  000001      No-override
  000010      1
  000011      No-override
  000100      No-override
  000101      5
  000110      No-override
  000111      7
  001000      No-override
     :             :
```

Priorities Configured in this step.

Figure 15-24. Assign Priorities to the Selected DSCPs

3. Assign the DSCP policies to the selected VIDs and display the result.

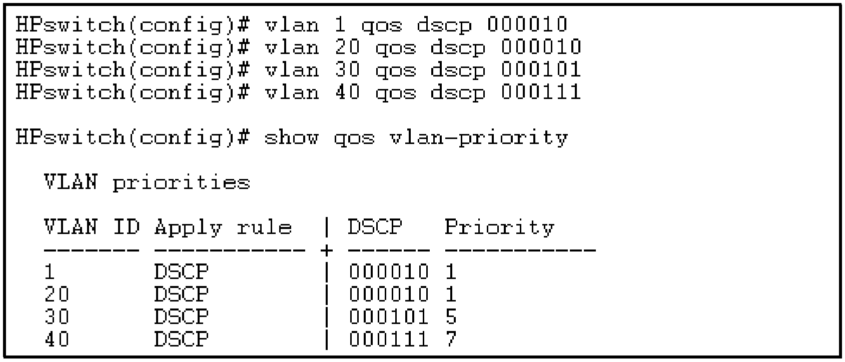

```
HPswitch(config)# vlan 1 qos dscp 000010
HPswitch(config)# vlan 20 qos dscp 000010
HPswitch(config)# vlan 30 qos dscp 000101
HPswitch(config)# vlan 40 qos dscp 000111

HPswitch(config)# show qos vlan-priority

  VLAN priorities

  VLAN ID Apply rule  | DSCP   Priority
  ------- ----------- + ------ -----------
  1       DSCP        | 000010 1
  20      DSCP        | 000010 1
  30      DSCP        | 000101 5
  40      DSCP        | 000111 7
```

Figure 15-25. The Completed VID-DSCP Priority Configuration

The switch will now apply the DSCP policies in figure 15-25 to packets received on the switch with the specified VLAN-IDs. This means the switch will:

- Overwrite the original DSCPs in the selected packets with the new DSCPs specified in the above policies.
- Assign the 802.1p priorities in the above policies to the appropriate packets.

QoS Source-Port Priority

QoS Classifier Precedence: 6

The QoS source-port option enables you to use a packet's source-port on the switch as a QoS classifier. Where a particular source-port classifier has the highest precedence in the switch for traffic entering through that port, then traffic received from the port is marked with the source-port classifier's configured priority level. Different source-port classifiers can have different priority levels.

Options for Assigning Priority. Priority control options for packets from a specified source-port include:

- 802.1p priority
- DSCP policy (Assigning a new DSCP and an associated 802.1p priority; inbound packets must be IPv4.)

(For operation when other QoS classifiers apply to the same traffic, refer to "Classifiers for Prioritizing Outbound Packets" on page 15-8.)

Assigning a Priority Based on Source-Port Only

This option assigns a priority to all outbound packets having the specified source-port. You can configure this option by either specifying the source-port ahead of the qos command or moving to the port context for the port you want to configure for priority. (If you are configuring multiple source-ports with the same priority, you may find it easier to use the port context instead of individually configuring the priority for each port.)

Syntax: interface [e] < *port-list* > qos priority < 0 - 7 >

Configures an 802.1p priority for outbound packets entering the switch through the specified (source) ports. This priority determines the packet queue in the outbound port(s) to which traffic from the source-ports is sent. If a packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. You can configure one QoS classifier for each source-port or group of source-ports. (Default: **No-override***)*

no interface [e] < *port-list* > qos

Disables use of the specified source-port(s) for QoS classifier(s) and resets the priority for the specified source-port(s) to **No-override**.

show qos vlan-priority

Lists the QoS VLAN-ID classifiers with their priority data.

For example, suppose that you want to prioritize inbound traffic on the following source-ports:

Source-Port	Priority
A1 - A3	2
A4	3
B1, B4	5
C1-C3	6

You would then execute the following commands to prioritize traffic received on the above ports:

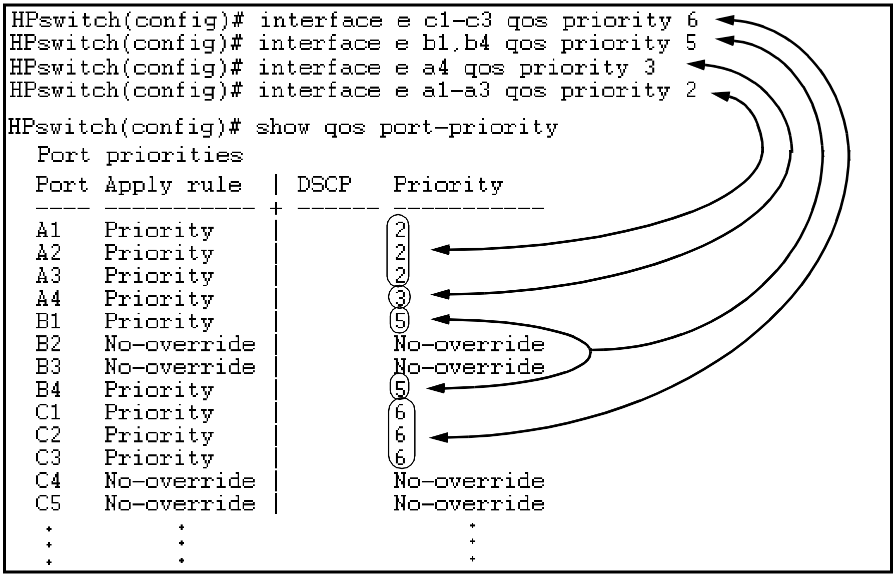

```
HPswitch(config)# interface e c1-c3 qos priority 6
HPswitch(config)# interface e b1,b4 qos priority 5
HPswitch(config)# interface e a4 qos priority 3
HPswitch(config)# interface e a1-a3 qos priority 2

HPswitch(config)# show qos port-priority
  Port priorities
  Port Apply rule  | DSCP   Priority
  ---- ----------- + ------ -----------
  A1   Priority    |        2
  A2   Priority    |        2
  A3   Priority    |        2
  A4   Priority    |        3
  B1   Priority    |        5
  B2   No-override |        No-override
  B3   No-override |        No-override
  B4   Priority    |        5
  C1   Priority    |        6
  C2   Priority    |        6
  C3   Priority    |        6
  C4   No-override |        No-override
  C5   No-override |        No-override
  .         .                   .
  .         .                   .
  .         .                   .
```

Figure 15-26. Configuring and Displaying Source-Port QoS Priorities

If you then decided to remove port A1 from QoS prioritization:

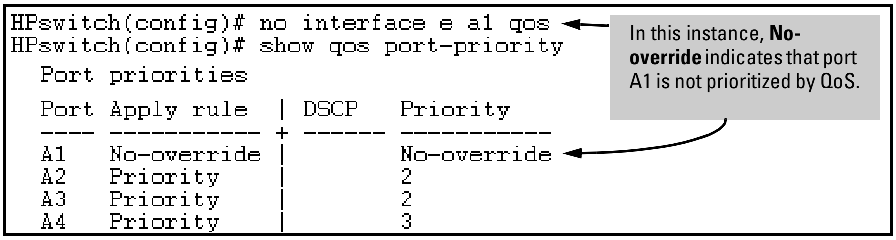

```
HPswitch(config)# no interface e a1 qos
HPswitch(config)# show qos port-priority
  Port priorities
  Port Apply rule  | DSCP   Priority
  ---- ----------- + ------ -----------
  A1   No-override |        No-override
  A2   Priority    |        2
  A3   Priority    |        2
  A4   Priority    |        3
```

In this instance, **No-override** indicates that port A1 is not prioritized by QoS.

Figure 15-27. Returning a QoS-Prioritized VLAN to "No-override" Status

Assigning a DSCP Policy Based on the Source-Port

This option assigns a previously configured DSCP policy (codepoint and 802.1p priority) to outbound IP packets (received from the specified source-ports). That is, the switch:

1. Selects an incoming IP packet on the basis of its source-port on the switch.
2. Overwrites the packet's DSCP with the DSCP configured in the switch for such packets.
3. Assigns the 802.1p priority configured in the switch for the new DSCP. (Refer to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)
4. Forwards the packet through the appropriate outbound port queue.

For more on DSCP, refer to "Terminology" on page 15-5.

Steps for Creating a Policy Based on Source-Port Classifiers.

Note

You can select one DSCP per source-port. Also, configuring a new DSCP for a source-port automatically overwrites (replaces) any previous DSCP or 802.1p priority configuration for that port.)

1. Identify the source-port classifier to which you want to assign a DSCP policy.
2. Determine the DSCP policy for packets having the selected source-port:
   a. Determine the DSCP you want to assign to the selected packets. (This codepoint will be used to overwrite the DSCP carried in packets received through the source-port from upstream devices.)
   b. Determine the 802.1p priority you want to assign to the DSCP.
3. Configure the DSCP policy by using **qos dscp-map** to configure the priority for each codepoint. (For details, refer to the example later in this section and to "Differentiated Services Codepoint (DSCP) Mapping" on page 15-51.)

Note

A codepoint must have an 802.1p priority assignment (0 - 7) before you can configure that codepoint as a criteria for prioritizing packets by source-port. If a codepoint shows **No-override** in the **Priority** column of the DSCP Policy Table (**show qos dscp-map**), then you must assign a 0 - 7 priority before proceeding.

4. Configure the switch to assign the DSCP policy to packets from the specified source-port.

Syntax: qos dscp-map < *codepoint* > priority < 0 - 7 >

This command is optional if a priority has already been assigned to the < codepoint*>. The command creates a DSCP policy by assigning an 802.1p priority to a specific DSCP. When the switch applies this priority to a packet, the priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. (Default: For most codepoints,* **No-override**. *See figure 15-8 on page 15-52.)*

Syntax: interface [e] < *port-list* > qos dscp < *codepoint* >

Assigns a DSCP policy to packets from the specified source-port(s), and overwrites the DSCP in these packets with the assigned < ***codepoint*** *> value. This policy includes an 802.1p priority and determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority with it to the next downstream device. (Default:* **No-override***)*

no interface [e] < *port-list* > qos

Removes QoS classifier for the specified source-port(s).

show qos source-port

Displays a listing of all source-port QoS classifiers currently in the running-config file.

For example, suppose you wanted to assign this set of priorities:

Source-Port	DSCP	Priority
A2	000111	7
B1-B3	000101	5
B4, C2	000010	1

1. Determine whether the DSCPs already have priority assignments, which could indicate use by existing applications. This is not a problem as long as the configured priorities are acceptable for all applications using the same DSCP. (Refer to the “Note On Changing a Priority Setting” on page 15-54. Also, a DSCP must have a priority configured before you can assign any QoS classifiers to use it.)

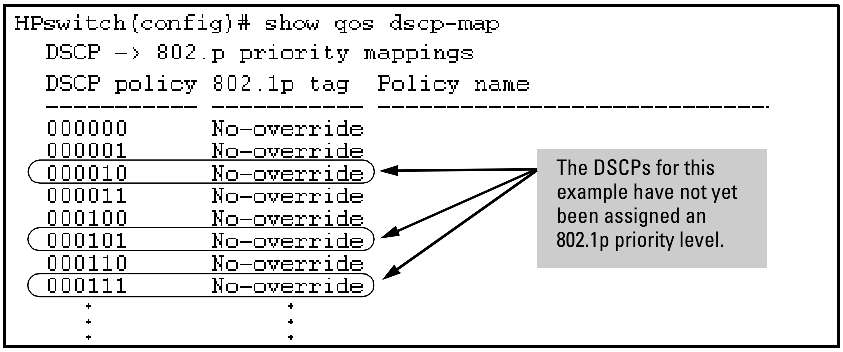

```
HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag  Policy name
  ----------- ----------- -------------------------------
  000000      No-override
  000001      No-override
  000010      No-override
  000011      No-override
  000100      No-override
  000101      No-override
  000110      No-override
  000111      No-override
     .           .
     .           .
     .           .
```

The DSCPs for this example have not yet been assigned an 802.1p priority level.

Figure 15-28. Display the Current Configuration in the DSCP Policy Table

2. Configure the priorities for the DSCPs you want to use.

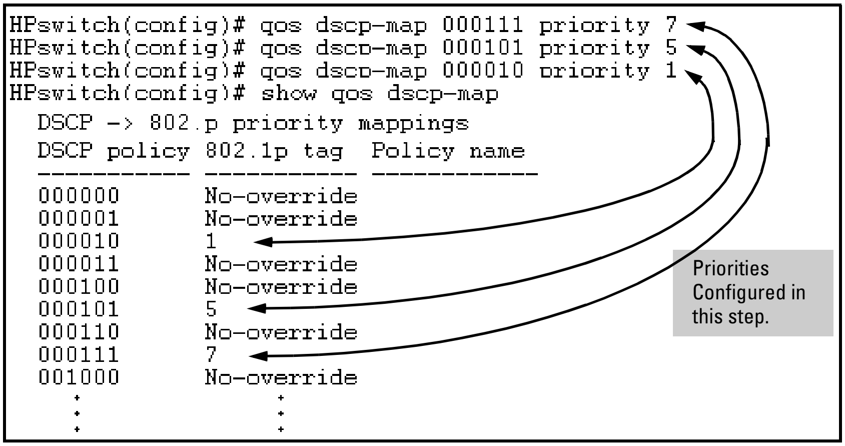

```
HPswitch(config)# qos dscp-map 000111 priority 7
HPswitch(config)# qos dscp-map 000101 priority 5
HPswitch(config)# qos dscp-map 000010 priority 1
HPswitch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag  Policy name
  ----------- ----------- ------------
  000000      No-override
  000001      No-override
  000010      1
  000011      No-override
  000100      No-override
  000101      5
  000110      No-override
  000111      7
  001000      No-override
     .           .
     .           .
     .           .
```

Priorities Configured in this step.

Figure 15-29. Assign Priorities to the Selected DSCPs

3. Assign the DSCP policies to the selected source-ports and display the result.

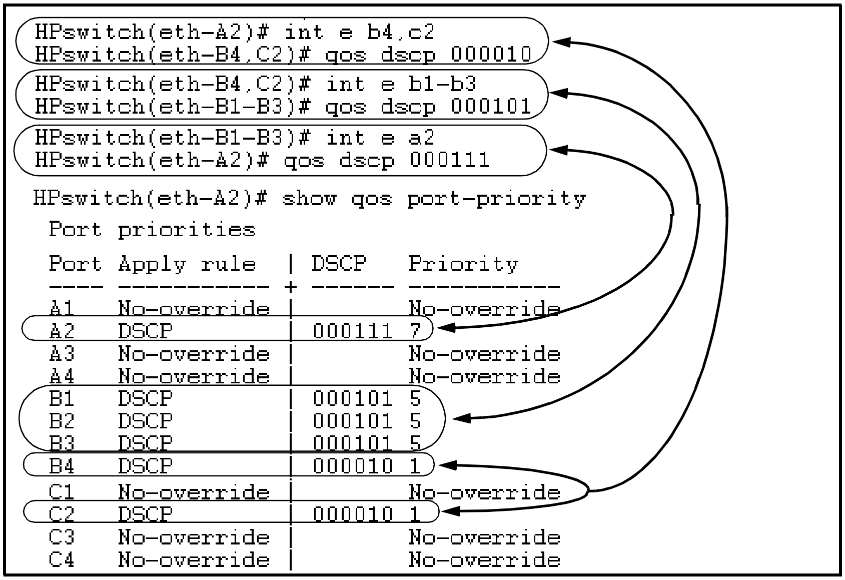

```
HPswitch(eth-A2)# int e b4,c2
HPswitch(eth-B4,C2)# qos dscp 000010
HPswitch(eth-B4,C2)# int e b1-b3
HPswitch(eth-B1-B3)# qos dscp 000101
HPswitch(eth-B1-B3)# int e a2
HPswitch(eth-A2)# qos dscp 000111

HPswitch(eth-A2)# show qos port-priority
 Port priorities

 Port Apply rule   | DSCP   Priority
 ---- ----------- + ------ -----------
 A1   No-override |        No-override
 A2   DSCP        | 000111 7
 A3   No-override |        No-override
 A4   No-override |        No-override
 B1   DSCP        | 000101 5
 B2   DSCP        | 000101 5
 B3   DSCP        | 000101 5
 B4   DSCP        | 000010 1
 C1   No-override |        No-override
 C2   DSCP        | 000010 1
 C3   No-override |        No-override
 C4   No-override |        No-override
```

Figure 15-30. The Completed Source-Port DSCP-Priority Configuration

Differentiated Services Codepoint (DSCP) Mapping

The DSCP Policy Table associates an 802.1p priority with a specific ToS byte codepoint in an IPv4 packet. This enables you to set a LAN policy that operates independently of 802.1Q VLAN-tagging.

In the default state, most of the 64 codepoints do not assign an 802.1p priority, as indicated by **No-override** in table 15-8 on page 15-52.

You can list the current DSCP Policy table, change the codepoint priority assignments, and assign optional names to the codepoints.

Syntax: show qos dscp-map

Displays the DSCP Policy Table.

qos dscp-map < *codepoint* > priority < 0 - 7 > [name < *ascii-string* >]

Configures an 802.1p priority for the specified codepoint and, optionally, an identifying (policy) name.

no qos dscp-map < *codepoint* >

Reconfigures the 802.1p priority for <codepoint> to **No-override**. *Also deletes the codepoint policy name, if configured.*

no qos dscp-map < *codepoint* > name

Deletes only the policy name, if configured, for < codepoint >.

Table 16-8. The Default DSCP Policy Table

DSCP Policy	802.1p Priority	DSCP Policy	802.1p Priority	DSCP Policy	802.1p Priority
000000	No-override	010110	3	101011	No-override
000001	No-override	010111	No-override	101100	No-override
000010	No-override	011000	No-override	101101	No-override
000011	No-override	011001	No-override	101110	7
000100	No-override	011010	4	101111	No-override
000101	No-override	011011	No-override	110000	No-override
000110	No-override	011100	4	110001	No-override
000111	No-override	011101	No-override	110010	No-override
001000	No-override	011110	5	110011	No-override
001001	No-override	011111	No-override	110100	No-override
001010	1	100000	No-override	110101	No-override
001011	No-override	100001	No-override	110110	No-override
001100	1	100010	6	110111	No-override
001101	No-override	100011	No-override	111000	No-override
001110	2	100100	6	111001	No-override
001111	No-override	100101	No-override	111010	No-override
010000	No-override	100110	7	111011	No-override
010001	No-override	100111	No-override	111100	No-override
010010	0	101000	No-override	111101	No-override
010011	No-override	101001	No-override	111110	No-override
010100	0	101010	No-override	111111	No-override
010101	No-override				

Default Priority Settings for Selected Codepoints

In a few cases, such as 001010 and 001100, a default policy (implied by the DSCP standards for Assured-Forwarding and Expedited-Forwarding) is used. You can change the priorities for the default policies by using **qos dscp-map < *codepoint* > priority < 0 - 7 >)**.(These policies are not in effect unless you have either applied the policies to a QoS classifier or configured QoS Type-of-Service to be in **diff-services** mode.)

Quickly Listing Non-Default Codepoint Settings

Table 15-8 lists the switch's default codepoint/priority settings. If you change the priority of any codepoint setting to a non-default value and then execute **write memory**, the switch will list the non-default setting in the show config display. For example, in the default configuration, the following codepoint settings are true:

Codepoint	Default Priority
001100	1
001101	No-override
001110	2

If you change all three settings to a priority of 3, and then execute **write memory**, the switch will reflect these changes in the show config listing:

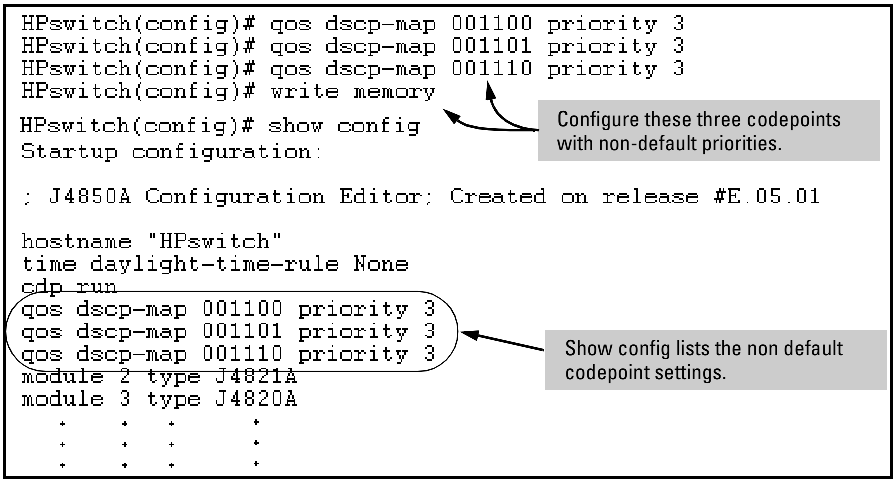

```
HPswitch(config)# qos dscp-map 001100 priority 3
HPswitch(config)# qos dscp-map 001101 priority 3
HPswitch(config)# qos dscp-map 001110 priority 3
HPswitch(config)# write memory
HPswitch(config)# show config
Startup configuration:

; J4850A Configuration Editor; Created on release #E.05.01

hostname "HPswitch"
time daylight-time-rule None
cdp run
qos dscp-map 001100 priority 3
qos dscp-map 001101 priority 3
qos dscp-map 001110 priority 3
module 2 type J4821A
module 3 type J4820A
   .    .   .      .
   .    .   .      .
   .    .   .      .
```

Configure these three codepoints with non-default priorities.

Show config lists the non default codepoint settings.

Figure 15-31. Example of Show Config Listing with Non-Default Priority Settings in the DSCP Table

Effect of "No-override". In the QoS Type-of-Service differentiated services mode, a **No-override** assignment for the codepoint of an outbound packet means that QoS is effectively disabled for such packets. That is, QoS does not affect the packet queuing priority or VLAN tagging. In this case, the packets are handled as follows (as long as no other QoS feature creates priority assignments for them):

802.1Q Status	Outbound 802.1p Priority
Received and Forwarded on a tagged VLAN	Unchanged
Received on an Untagged VLAN; Forwarded on a tagged VLAN	0 (zero)—"normal"
Forwarded on an Untagged VLAN	None

Note On Changing a Priority Setting

If a QoS classifier is using a policy (codepoint and associated priority) in the DSCP Policy table, you must delete or change this usage before you can change the priority setting on the codepoint. Otherwise the switch blocks the change and displays this message:

> **Cannot modify DSCP Policy < *codepoint* > - in use by other qos rules.**

In this case, use **show qos < *classifier* >** to identify the specific classifiers using the policy you want to change; that is:

> **show qos device-priority**
> **show qos port-priority**
> **show qos tcp-udp-port-priority**
> **show qos vlan-priority**
> **show qos type-of-service**

Note that protocol-priority is not included because a DSCP policy is not meaningful for this classifier and therefore not configurable in this case.

For example, suppose that the 000001 codepoint has a priority of 6, and several classifiers use the 000001 codepoint to assign a priority to their respective types of traffic. If you wanted to change the priority of codepoint 000001 you would do the following:

1. Identify which QoS classifiers use the codepoint.
2. Change the classifier configurations by assigning them to a different DSCP policy, or to an 802.1p priority, or to **No-override**.
3. Reconfigure the desired priority for the 000001 codepoint.
4. Either reassign the classifiers to the 00001 codepoint policy or leave them as they were after step 2, above.

Example of Changing the Priority Setting on a Policy When One or More Classifiers Are Currently Using the Policy

Suppose that codepoint 000001 is in use by one or more classifiers. If you try to change its priority, you see a result similar to the following:

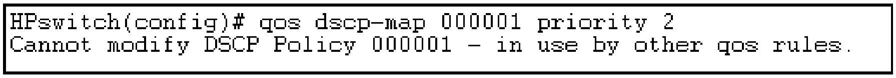

```
HPswitch(config)# qos dscp-map 000001 priority 2
Cannot modify DSCP Policy 000001 - in use by other qos rules.
```

Figure 15-32. Example of Trying To Change the Priority on a Policy In Use by a Classifier

In this case, you would use steps similar to the following to change the priority.

1. Identify which classifiers use the codepoint you want to change.

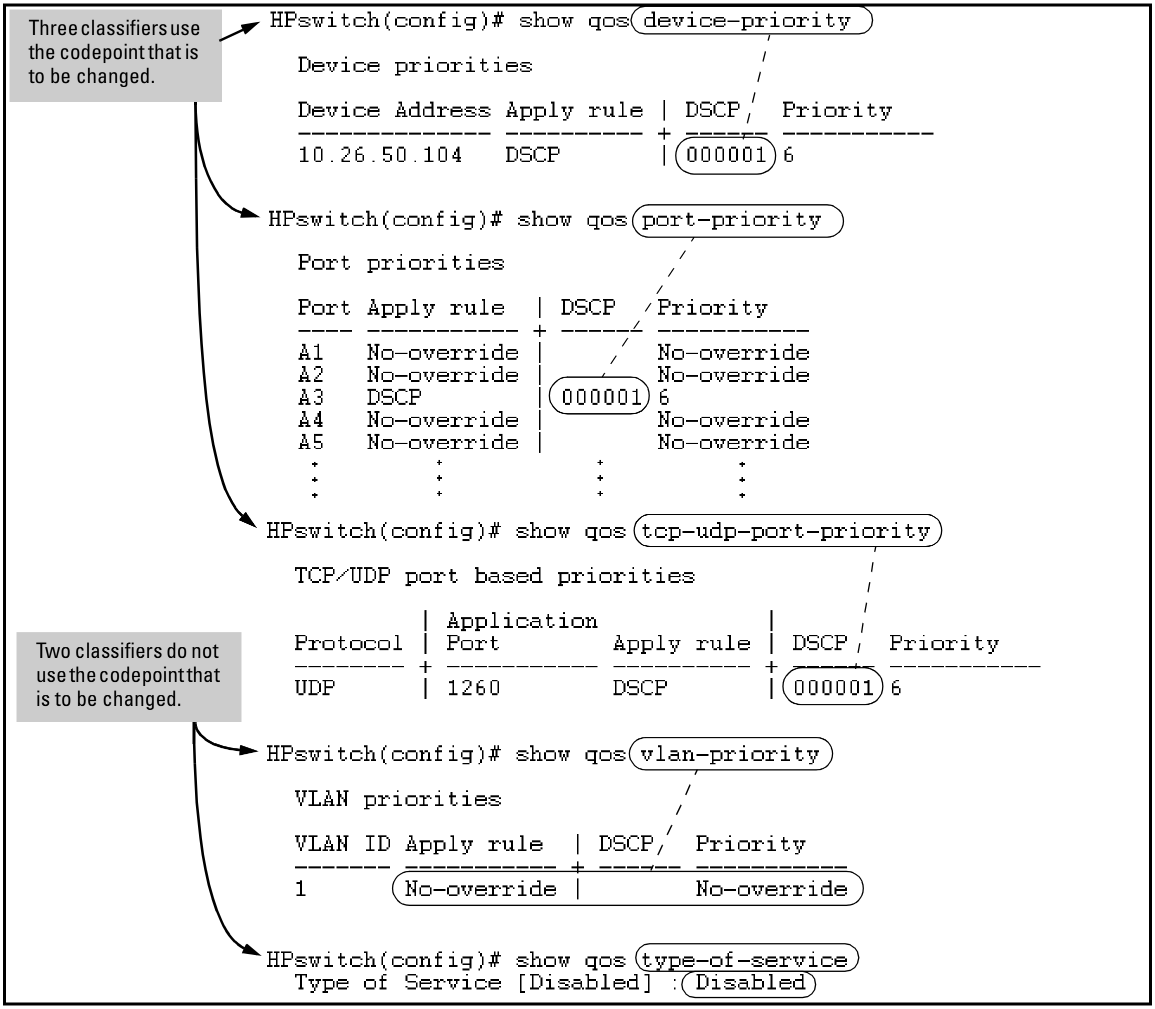

Three classifiers use the codepoint that is to be changed.

Two classifiers do not use the codepoint that is to be changed.

```
HPswitch(config)# show qos device-priority

  Device priorities

  Device Address Apply rule | DSCP   Priority
  -------------- ---------- + ------ -----------
  10.26.50.104   DSCP       | 000001 6

HPswitch(config)# show qos port-priority

  Port priorities

  Port Apply rule  | DSCP   Priority
  ---- ----------- + ------ -----------
  A1   No-override |        No-override
  A2   No-override |        No-override
  A3   DSCP        | 000001 6
  A4   No-override |        No-override
  A5   No-override |        No-override
   .        .           .         .
   .        .           .         .
   .        .           .         .
HPswitch(config)# show qos tcp-udp-port-priority

  TCP/UDP port based priorities

           | Application                       |
  Protocol | Port        Apply rule | DSCP   Priority
  -------- + ----------- ---------- + ------ -----------
  UDP      | 1260        DSCP       | 000001 6

HPswitch(config)# show qos vlan-priority

  VLAN priorities

  VLAN ID Apply rule  | DSCP   Priority
  ------- ----------- + ------ -----------
  1       No-override |        No-override

HPswitch(config)# show qos type-of-service
  Type of Service [Disabled] : Disabled
```

Figure 15-33. Example of a Search to Identify Classifiers Using a Codepoint You Want To Change

2. Change the classifier configurations by assigning them to a different DSCP policy, or to an 802.1p priority, or to **No-override**. For example:
   a. Delete the policy assignment for the **device-priority** classifier. (That is, assign it to **No-override**.)
   b. Create a new DSCP policy to use for re-assigning the remaining classifiers.
   c. Assign the **port-priority** classifier to the new DSCP policy.
   d. Assign the **udp-port 1260** classifier to an 802.1p priority.

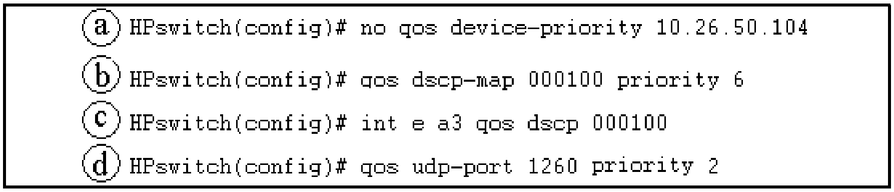

```
HPswitch(config)# no qos device-priority 10.26.50.104
HPswitch(config)# qos dscp-map 000100 priority 6
HPswitch(config)# int e a3 qos dscp 000100
HPswitch(config)# qos udp-port 1260 priority 2
```

3. Reconfigure the desired priority for the 000001 codepoint.

```
HPswitch(config)# qos dscp-map 000001 priority 4k
```

4. You could now re-assign the classifiers to the original policy codepoint or leave them as currently configured.

Configuring QoS from the Web Browser Interface

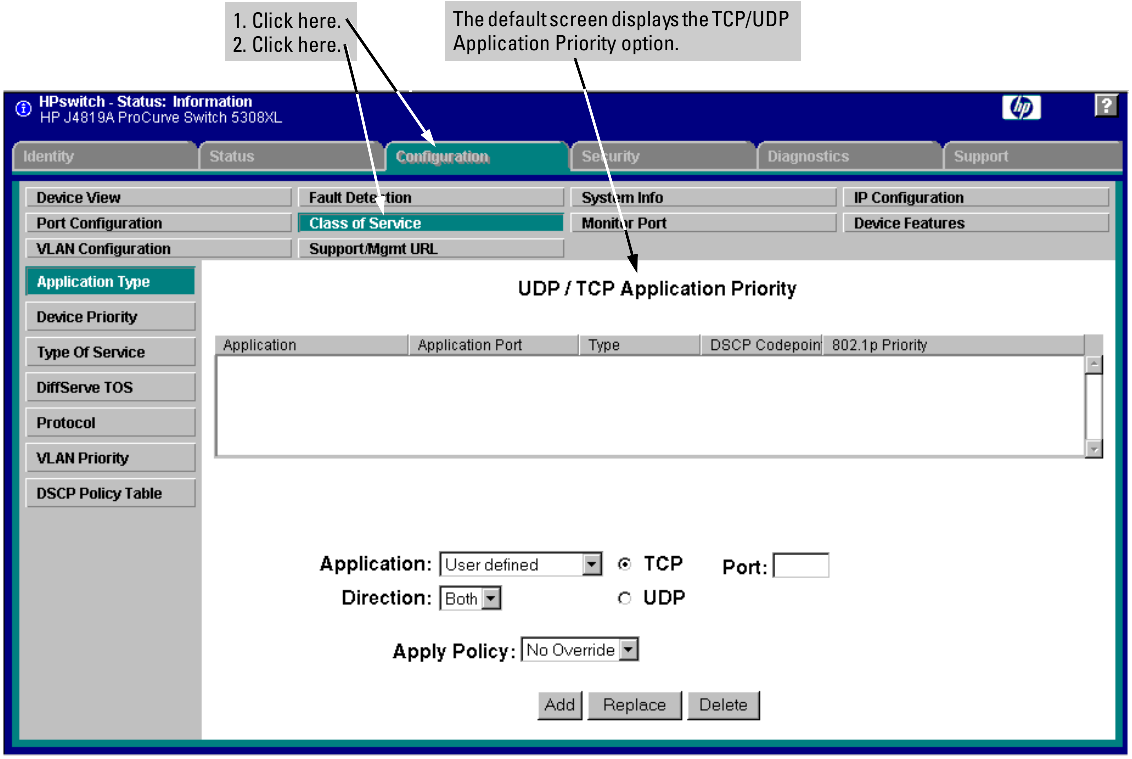

1. Click here.
2. Click here.

The default screen displays the TCP/UDP Application Priority option.

HPswitch - Status: Information
HP J4819A ProCurve Switch 5308XL

UDP / TCP Application Priority

Application	Application Port	Type	DSCP Codepoin	802.1p Priority

Application: User defined ● TCP Port:
Direction: Both ○ UDP

Apply Policy: No Override

Figure 15-34. The Default QoS Configuration Screen

IP Multicast (IGMP) Interaction with QoS

IGMP high-priority-forward causes the switch to service the subscribed IP multicast group traffic at high priority, even if QoS on the switch has relegated the traffic to a lower priority. This does not affect any QoS priority settings, so the QoS priority is honored by downstream devices. However, QoS does take precedence over IGMP normal-priority traffic.

The switch's ability to prioritize IGMP traffic for either a normal or high priority outbound queue overrides any QoS criteria, and does not affect any 802.1p priority settings the switch may assign. For a given packet, if both IGMP high priority and QoS are configured, the QoS classification occurs and the switch marks the packet for downstream devices, but the packet is serviced by the high-priority queue when leaving the switch.

IGMP High Priority	QoS Configuration Affects Packet	Switch Port Output Queue	Outbound 802.1p Setting (Requires Tagged VLAN)
Not Enabled	Yes	Determined by QoS	Determined by QoS
Enabled	See above paragraph.	High	As determined by QoS if QoS is active.

QoS Messages in the CLI

Message	Meaning
DSCP Policy < *decimal-codepoint* > not configured	You have attempted to map a QoS classifier to a codepoint for which there is no configured priority (**No-override**). Use the **qos dscp-map** command to configure a priority for the codepoint, then map the classifier to the codepoint.
Cannot modify DSCP Policy < *codepoint* > - in use by other qos rules.	You have attempted to map a QoS classifier to a codepoint that is already in use by other QoS classifiers. Before remapping the codepoint to a new priority, you must reconfigure the other QoS classifiers so that they do not use this codepoint. You can have multiple QoS classifiers use this same codepoint as long as it is acceptable for all such classifiers to use the same priority.

Operating and Troubleshooting Notes

- **For Devices that Do Not Support 802.1Q Tagged VLANs:** For communication between these devices and the switch, connect the device to a switch port configured as **Untagged** for the VLAN in which you want the device's traffic to move.

- **VLAN Tagging Rules:** For a port on the switch to be a member of a VLAN, the port must be configured as either **Tagged** or **Untagged** for that VLAN. (Only one VLAN on a port can be untagged. Otherwise, the switch cannot determine which VLAN should receive untagged VLAN traffic.)

- **Loss of Communication on a Tagged VLAN:** If you cannot communicate with a device in a tagged VLAN environment, ensure that the device either supports tagged VLANs or is connected to a VLAN port that is configured as **Untagged**.

16

IP Routing Features

Contents

Overview of IP Routing

The HP Procurve Series 5300XL Switches offer the following IP routing features:

- **IP Static Routes** – supports up to 256 static routes
- **RIP** (Router Information Protocol) – supports RIP Version 1, Version 1 compatible with Version 2 (default), and Version 2
- **OSPF** (Open Shortest Path First) – the standard routing protocol for handling larger routed networks
- **IRDP** (ICMP Router Discovery Protocol) – advertises the IP addresses of the routing interfaces on this switch to directly attached host systems
- **DHCP Relay** – allows you to extend the service range of your DHCP server beyond its single local network segment

All these features are configurable through the switch's console CLI.

Throughout this chapter, the HP Procurve Series 5300XL Switches will be referred to as "routing switches". When IP routing is enabled on your switch, it behaves just like any other IP router.

Basic IP routing configuration consists of adding IP addresses, enabling IP routing, and, enabling a route exchange protocol, such as Routing Information Protocol (RIP).

For configuring the IP addresses, see chapter 7, "Configuring IP Addresses". The rest of this chapter describes IP routing and how to configure it in more detail. Use the information in this chapter if you need to change some of the IP parameters from their default values or you want to view configuration information or statistics.

IP Interfaces

On the routing switches, IP addresses are associated with individual VLANs. By default, there is a single VLAN (Default_VLAN) on the routing switch. In that configuration, a single IP address serves as the management access address for the entire device. If routing is enabled on the routing switch, the IP address on the single VLAN also acts as the routing interface.

Each IP address on a routing switch must be in a different sub-net. You can have only one VLAN interface that is in a given sub-net. For example, you can configure IP addresses 192.168.1.1/24 and 192.168.2.1/24 on the same routing switch, but you cannot configure 192.168.1.1/24 and 192.168.1.2/24 on the same routing switch.

You can configure multiple IP addresses on the same VLAN.

The number of IP addresses you can configure on an individual VLAN interface is 8.

You can use any of the IP addresses you configure on the routing switch for Telnet, Web management, or SNMP access, as well as for routing.

Note

All HP Procurve devices support configuration and display of IP address in classical sub-net format (example: 192.168.1.1 255.255.255.0) and Classless Interdomain Routing (CIDR) format (example: 192.168.1.1/24). You can use either format when configuring IP address information. IP addresses are displayed in classical sub-net format only.

IP Tables and Caches

The following sections describe the IP tables and caches:

- ARP cache table
- IP route table
- IP forwarding cache

The software enables you to display these tables.

ARP Cache Table

The ARP cache contains entries that map IP addresses to MAC addresses. Generally, the entries are for devices that are directly attached to the routing switch.

An exception is an ARP entry for an interface-based static IP route that goes to a destination that is one or more router hops away. For this type of entry, the MAC address is either the destination device's MAC address or the MAC address of the router interface that answered an ARP request on behalf of the device, using proxy ARP.

ARP Cache. The ARP cache contains dynamic (learned) entries. The software places a dynamic entry in the ARP cache when the routing switch learns a device's MAC address from an ARP request or ARP reply from the device.

The software can learn an entry when the switch or routing switch receives an ARP request from another IP forwarding device or an ARP reply. Here is an example of a dynamic entry:

```
     IP Address         MAC Address        Type         Port
1     207.95.6.102      0800.5afc.ea21    Dynamic       6
```

Each entry contains the destination device's IP address and MAC address.

To configure other ARP parameters, see "Configuring ARP Parameters" on page 16-10.

IP Route Table

The IP route table contains routing paths to IP destinations.

Note

The default gateway, which you specify when you configure the basic IP information on the switch, is used only when routing is not enabled on the switch.

The IP route table can receive the routing paths from the following sources:

- A directly-connected destination, which means there are no router hops to the destination
- A static IP route, which is a user-configured route
- A route learned through RIP
- A route learned through OSPF

The IP route table contains the best path to a destination.

- When the software receives paths from more than one of the sources listed above, the software compares the administrative distance of each path and selects the path with the lowest administrative distance. The administrative distance is a protocol-independent value from 1 – 255.

The IP route table is displayed by entering the CLI command **show ip route** from any context level in the console CLI. Here is an example of an entry in the IP route table:

```
Destination    Network Mask     Gateway    Type         Sub-Type     Metric
1.1.0.0       255.255.0.0     99.1.1.2   connected                  1
```

Each IP route table entry contains the destination's IP address and sub-net mask and the IP address of the next-hop router interface to the destination. Each entry also indicates route type, and for OSPF routes, the sub type, and the route's IP metric (cost). The type indicates how the IP route table received the route.

To configure a static IP route, see "Configuring a Static IP Route" on page 16-18

IP Forwarding Cache

The IP forwarding cache provides a fast-path mechanism for forwarding IP packets. The cache contains entries for IP destinations. When an HP ProCurve routing switch has completed processing and addressing for a packet and is ready to forward the packet, the device checks the IP forwarding cache for an entry to the packet's destination.

- If the cache contains an entry with the destination IP address, the device uses the information in the entry to forward the packet out the ports listed in the entry. The destination IP address is the address of the packet's final destination. The port numbers are the ports through which the destination can be reached.
- If the cache does not contain an entry, the software can create an entry in the forwarding cache.

Each entry in the IP forwarding cache has an age timer. If the entry remains unused for five minutes, the software removes the entry. The age timer is not configurable.

Note

You cannot add static entries to the IP forwarding cache.

IP Route Exchange Protocols

HP Procurve Series 5300XL Switches support the following IP route exchange protocols:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)

These protocols provide routes to the IP route table. You can use one or more of these protocols, in any combination. The protocols are disabled by default. For configuration information, see the following:

- "Configuring RIP" on page 16-20
- "Configuring OSPF" on page 16-33

IP Global Parameters for Routing Switches

The following table lists the IP global parameters and the page where you can find more information about each parameter.

Table 16-1. IP Global Parameters for Routing Switches

Parameter	Description	Default	See page
Router ID	The value that routers use to identify themselves to other routers when exchanging route information. OSPF uses the router ID to identify routers. RIP does not use the router ID.	The lowest-numbered IP address configured on the lowest-numbered routing interface.	16-9
Address Resolution Protocol (ARP)	A standard IP mechanism that routers use to learn the Media Access Control (MAC) address of a device on the network. The router sends the IP address of a device in the ARP request and receives the device's MAC address in an ARP reply.	Enabled	16-10
ARP age	The amount of time the device keeps a MAC address learned through ARP in the device's ARP cache. The device resets the timer to zero each time the ARP entry is refreshed and removes the entry if the timer reaches the ARP age.	Five minutes	not configurable
Proxy ARP	An IP mechanism a router can use to answer an ARP request on behalf of a host, by replying with the router's own MAC address instead of the host's.	Disabled	16-12
Time to Live (TTL)	The maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet's TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it.	64 hops	7-11

Parameter	Description	Default	See page
Directed broadcast forwarding	A directed broadcast is a packet containing all ones (or in some cases, all zeros) in the host portion of the destination IP address. When a router forwards such a broadcast, it sends a copy of the packet out each of its enabled IP interfaces. **Note**: You also can enable or disable this parameter on an individual interface basis. See table 16-2 on page 16-8.	Disabled	16-13
ICMP Router Discovery Protocol (IRDP)	An IP protocol that a router can use to advertise the IP addresses of its router interfaces to directly attached hosts. You can enable or disable the protocol at the Global CLI Config level. You also can enable or disable IRDP and configure the following protocol parameters on an individual VLAN interface basis at the VLAN Interface CLI Config level. • Forwarding method (broadcast or multicast) • Hold time • Maximum advertisement interval • Minimum advertisement interval • Router preference level	Disabled	16-69 16-70
Static route	An IP route you place in the IP route table.	No entries	16-16
Default network route	The router uses the default network route if the IP route table does not contain a route to the destination. For the Series 5300XL Switches, enter an explicit default route (0.0.0.0 0.0.0.0 or 0.0.0.0/0) as a static route in the IP route table.	None configured	16-19

IP Interface Parameters for Routing Switches

Table 16-2 lists the interface-level IP parameters for routing switches.

Table 16-2. IP Interface Parameters – Routing Switches

Parameter	Description	Default	See page
IP address	A Layer 3 network interface address; separate IP addresses on individual VLAN interfaces.	None configured	chapter 7
Metric	A numeric cost the router adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1 (one)	16-22
ICMP Router Discovery Protocol (IRDP)	Locally overrides the global IRDP settings. See table 16-1 on page 16-6 for global IRDP information.	Disabled	16-70
IP helper address	The IP address of a UDP application server (such as a BootP or DHCP server) or a directed broadcast address. IP helper addresses allow the routing switch to forward requests for certain UDP applications from a client on one sub-net to a server on another sub-net.	None configured	16-73

Configuring IP Parameters for Routing Switches

The following sections describe how to configure IP parameters. Some parameters can be configured globally while others can be configured on individual VLAN interfaces. Some parameters can be configured globally and overridden for individual VLAN interfaces.

Note

This section describes how to configure IP parameters for routing switches. For IP configuration information when routing is not enabled, see chapter 7, “Configuring IP Addressing”.

Configuring IP Addresses

You can configure an IP address on the routing switch’s VLAN interfaces. Configuring IP addresses is described in detail in chapter 7, “Configuring IP Addressing”.

Changing the Router ID

In most configurations, a routing switch has multiple IP addresses, usually configured on different VLAN interfaces. As a result, a routing switch’s identity to other devices varies depending on the interface to which the other device is attached. Some routing protocols, including Open Shortest Path First (OSPF), identify a routing switch by just one of the IP addresses configured on the routing switch, regardless of the interfaces that connect the routing switches. This IP address is the router ID.

Note

Routing Information Protocol (RIP) does not use the router ID.

By default, the router ID on an HP routing switch is the lowest numbered IP interface configured on the device.

If you prefer, you can explicitly set the router ID to any valid IP address. The IP address cannot be in use on another device in the network.

Note

To display the router ID, enter the **show ip ospf** CLI command at any Manager EXEC CLI level.

To change the router ID, enter a command such as the following:

```
HPswitch(config)# ip router-id 209.157.22.26
```

Syntax: Syntax: ip router-id <ip-addr>

The **<ip-addr>** can be any valid, unique IP address.

Note

You can specify an IP address used for an interface on the HP routing switch, but do not specify an IP address in use by another device.

Configuring ARP Parameters

Address Resolution Protocol (ARP) is a standard IP protocol that enables an IP routing switch to obtain the MAC address of another device's interface when the routing switch knows the IP address of the interface. ARP is enabled by default and cannot be disabled.

How ARP Works

A routing switch needs to know a destination's MAC address when forwarding traffic, because the routing switch encapsulates the IP packet in a Layer 2 packet (MAC layer packet) and sends the Layer 2 packet to a MAC interface on a device directly attached to the routing switch. The device can be the packet's final destination or the next-hop router toward the destination.

The routing switch encapsulates IP packets in Layer 2 packets regardless of whether the ultimate destination is locally attached or is multiple router hops away. Since the routing switch's IP route table and IP forwarding cache contain IP address information but not MAC address information, the routing switch cannot forward IP packets based solely on the information in the route table or forwarding cache. The routing switch needs to know the MAC address that corresponds with the IP address of either the packet's locally attached destination or the next-hop router that leads to the destination.

For example, to forward a packet whose destination is multiple router hops away, the routing switch must send the packet to the next-hop router toward its destination, or to a default route or default network route if the IP route table does not contain a route to the packet's destination. In each case, the routing switch must encapsulate the packet and address it to the MAC address of a locally attached device, the next-hop router toward the IP packet's destination.

To obtain the MAC address required for forwarding a datagram, the routing switch does the following:

- First, the routing switch looks in the ARP cache (not the static ARP table) for an entry that lists the MAC address for the IP address. The ARP cache maps IP addresses to MAC addresses. The cache also lists the port attached to the device and, if the entry is dynamic, the age of the entry. A dynamic ARP entry enters the cache when the routing switch receives an ARP reply or receives an ARP request (which contains the sender's IP address and MAC address). A static entry enters the ARP cache from the static ARP table (which is a separate table) when the interface for the entry comes up.

  To ensure the accuracy of the ARP cache, each dynamic entry has its own age timer. The timer is reset to zero each time the routing switch receives an ARP reply or ARP request containing the IP address and MAC address of the entry. If a dynamic entry reaches its maximum allowable age, the entry times out and the software removes the entry from the table. Static entries do not age out and can be removed only by you.
- If the ARP cache does not contain an entry for the destination IP address, the routing switch broadcasts an ARP request out all its IP interfaces. The ARP request contains the IP address of the destination. If the device with the IP address is directly attached to the routing switch, the device sends an ARP response containing its MAC address. The response is a unicast packet addressed directly to the routing switch. The routing switch places the information from the ARP response into the ARP cache.

  ARP requests contain the IP address and MAC address of the sender, so all devices that receive the request learn the MAC address and IP address of the sender and can update their own ARP caches accordingly.

  Note: The ARP request broadcast is a MAC broadcast, which means the broadcast goes only to devices that are directly attached to the routing switch. A MAC broadcast is not routed to other networks. However, some routers, including HP routing switches, can be configured to reply to ARP requests from one network on behalf of devices on another network. See "Enabling Proxy ARP" below.

Note

If the routing switch receives an ARP request packet that it is unable to deliver to the final destination because of the ARP time-out and no ARP response is received (the routing switch knows of no route to the destination address), the routing switch sends an ICMP Host Unreachable message to the source.

Enabling Proxy ARP

Proxy ARP allows a routing switch to answer ARP requests from devices on one network on behalf of devices in another network. Since ARP requests are MAC-layer broadcasts, they reach only the devices that are directly connected to the sender of the ARP request. Thus, ARP requests do not cross routers.

For example, if Proxy ARP is enabled on a routing switch connected to two sub-nets, 10.10.10.0/24 and 20.20.20.0/24, the routing switch can respond to an ARP request from 10.10.10.69 for the MAC address of the device with IP address 20.20.20.69. In standard ARP, a request from a device in the 10.10.10.0/24 sub-net cannot reach a device in the 20.20.20.0 sub-net if the sub-nets are on different network cables, and thus is not answered.

An ARP request from one sub-net can reach another sub-net when both sub-nets are on the same physical segment (Ethernet cable), since MAC-layer broadcasts reach all the devices on the segment.

Proxy ARP is disabled by default on HP routing switches. To enable Proxy ARP, enter the following commands from the VLAN context level in the CLI:

```
HPswitch(config)# vlan 1
HPswitch(vlan-1)# ip proxy-arp
```

To again disable IP proxy ARP, enter the following command:

```
HPswitch(vlan-1)# no ip proxy-arp
```

Syntax: [no] ip proxy-arp

Configuring Forwarding Parameters

The following configurable parameters control the forwarding behavior of HP routing switches:

- Time-To-Live (TTL) threshold
- Forwarding of directed broadcasts

All these parameters are global and thus affect all IP interfaces configured on the routing switch.

To configure these parameters, use the procedures in the following sections.

Changing the TTL Threshold

The configuration of this parameter is covered in chapter 7, “Configuring IP Addressing”.

Enabling Forwarding of Directed Broadcasts

A directed broadcast is an IP broadcast to all devices within a single directly-attached network or sub-net. A net-directed broadcast goes to all devices on a given network. A sub-net-directed broadcast goes to all devices within a given sub-net.

Note

A less common type, the all-sub-nets broadcast, goes to all directly-attached sub-nets. Forwarding for this broadcast type also is supported, but most networks use IP multicasting instead of all-sub-net broadcasting.

Forwarding for all types of IP directed broadcasts is disabled by default. You can enable forwarding for all types if needed. You cannot enable forwarding for specific broadcast types.

To enable forwarding of IP directed broadcasts, enter the following CLI command:

```
HPswitch(config)# ip directed-broadcast
```

Syntax: [no] ip directed-broadcast

HP software makes the forwarding decision based on the routing switch's knowledge of the destination network prefix. Routers cannot determine that a message is unicast or directed broadcast apart from the destination network prefix. The decision to forward or not forward the message is by definition only possible in the last hop router.

To disable the directed broadcasts, enter the following CLI command:

```
HPswitch(config)# no ip directed-broadcast
```

Configuring ICMP

You can configure the following ICMP limits:

- **Burst-Normal** – The maximum number of ICMP replies to send per second.
- **Reply Limit** – You can enable or disable ICMP reply rate limiting.

Disabling ICMP Messages

HP devices are enabled to reply to ICMP echo messages and send ICMP Destination Unreachable messages by default.

You can selectively disable the following types of Internet Control Message Protocol (ICMP) messages:

- **Echo messages** (ping messages) – The routing switch replies to IP pings from other IP devices.
- **Destination Unreachable messages** – If the routing switch receives an IP packet that it cannot deliver to its destination, the routing switch discards the packet and sends a message back to the device that sent the packet to the routing switch. The message informs the device that the destination cannot be reached by the routing switch.
- **Address Mask replies** – You can enable or disable ICMP address mask replies.

Disabling Replies to Broadcast Ping Requests

By default, HP devices are enabled to respond to broadcast ICMP echo packets, which are ping requests. You can disable response to ping requests on a global basis using the following CLI method.

To disable response to broadcast ICMP echo packets (ping requests), enter the following command:

```
HPswitch(config)# no ip icmp echo broadcast-request
```

Syntax: [no] ip icmp echo broadcast-request

If you need to re-enable response to ping requests, enter the following command:

```
HPswitch(config)# ip icmp echo broadcast-request
```

Disabling ICMP Destination Unreachable Messages

By default, when an HP device receives an IP packet that the device cannot deliver, the device sends an ICMP Unreachable message back to the host that sent the packet. The following types of ICMP Unreachable messages are generated:

- Administration – The packet was dropped by the HP device due to a filter or ACL configured on the device.
- Fragmentation-needed – The packet has the Don't Fragment bit set in the IP Flag field, but the HP device cannot forward the packet without fragmenting it.
- Host – The destination network or sub-net of the packet is directly connected to the HP device, but the host specified in the destination IP address of the packet is not on the network.
- Network – The HP device cannot reach the network specified in the destination IP address of the packet.
- Port – The destination host does not have the destination TCP or UDP port specified in the packet. In this case, the host sends the ICMP Port Unreachable message to the HP device, which in turn sends the message to the host that sent the packet.
- Protocol – The TCP or UDP protocol on the destination host is not running. This message is different from the Port Unreachable message, which indicates that the protocol is running on the host but the requested protocol port is unavailable.
- Source-route-failure – The device received a source-routed packet but cannot locate the next-hop IP address indicated in the packet's Source-Route option.

Note

Disabling an ICMP Unreachable message type does not change the HP device's ability to forward packets. Disabling ICMP Unreachable messages prevents the device from generating or forwarding the Unreachable messages.

To disable all ICMP Unreachable messages, enter the following command:

```
HPswitch(config)# no ip icmp unreachable
```

Syntax: [no] ip icmp unreachable

Disabling ICMP Redirects

You can disable ICMP redirects on the HP routing switch. only on a global basis, for all the routing switch interfaces. To disable ICMP redirects globally, enter the following command at the global CONFIG level of the CLI:

```
HPswitch(config)# no ip icmp redirects
```

Syntax: [no] ip icmp redirects

Configuring Static IP Routes

The IP route table can receive routes from the following sources:

- Directly-connected networks – When you add an IP VLAN interface, the routing switch automatically creates a route for the network the interface is in.
- RIP – If RIP is enabled, the routing switch can learn about routes from the advertisements other RIP routers send to the routing switch. If the route has a lower administrative distance than any other routes from different sources to the same destination, the routing switch places the route in the IP route table.
- OSPF – See RIP, but substitute “OSPF” for “RIP”.
- Statically configured route – You can add routes directly to the route table. When you add a route to the IP route table, you are creating a static IP route. This section describes how to add static routes to the IP route table.
- Default network route – This is a specific static route that the routing switch uses if other routes to the destination are not available. See “Configuring the Default Route” on page 16-19.

Static Route Types

You can configure the following types of static IP routes:

- **Standard** – the static route consists of the destination network address and network mask, and the IP address of the next-hop gateway. You can configure multiple standard static routes with the same metric for load sharing or with different metrics to provide a primary route and backup routes.

- **Null (reject)** – the static route consists of the destination network address and network mask, and the **reject** parameter. Typically, the null route is configured as a backup route for discarding traffic if the primary route is unavailable.

Static IP Route Parameters

When you configure a static IP route, you must specify the following parameters:

- The IP address and network mask for the route's destination network.
- The route's path, which can be one of the following:
  - The IP address of a next-hop gateway
  - A "null" interface. The routing switch drops traffic forwarded to the null interface.

You also can specify the following optional parameters:

- The route's metric – The value the routing switch uses when comparing this route to other routes in the IP route table to the same destination. The metric applies only to routes that the routing switch has already placed in the IP route table. The default metric for static IP routes is 1.
- The route's administrative distance – The value that the routing switch uses to compare this route with routes from other route sources to the same destination before placing a route in the IP route table. This parameter does not apply to routes that are already in the IP route table. The default administrative distance for static IP routes is 1.

The default metric and administrative distance values ensure that the routing switch always prefers static IP routes over routes from other sources to the same destination.

Static Route States Follow Port States

IP static routes remain in the IP route table only so long as the next-hop gateway, port, or virtual interface used by the route is available. If the gateway or port becomes unavailable, the software removes the static route from the IP route table. If the gateway or port later becomes available again, the software adds the route back to the route table.

This feature allows the routing switch to adjust to changes in network topology. The routing switch does not continue trying to use routes on unavailable paths but instead uses routes only when their paths are available.

The following command configures a static route to 207.95.7.0, using 207.95.6.157 as the next-hop gateway.

```
HPswitch(config)# ip route 207.95.7.0/24 207.95.6.157
```

When you configure a static IP route, you specify the destination address for the route and the next-hop gateway or routing switch interface through which the routing switch can reach the route. The routing switch adds the route to the IP route table. In this case, Router A knows that 207.95.6.157 is reachable through port A2, and also assumes that local interfaces within that sub-net are on the same port. Router A deduces that IP interface 207.95.7.188 is also on port A2.

The software automatically removes a static IP route from the IP route table if the port used by that route becomes unavailable. When the port becomes available again, the software automatically re-adds the route to the IP route table.

Configuring a Static IP Route

To configure an IP static route with a destination address of 192.0.0.0 255.0.0.0 and a next-hop router IP address of 195.1.1.1, enter the following commands:

```
HPswitch(config)# ip route 192.0.0.0 255.0.0.0 195.1.1.1
```

Syntax: ip route <dest-ip-addr>/<mask-bits> <next-hop-ip-addr>

The <dest-ip-addr> is the route's destination. The <dest-mask> is the network mask for the route's destination IP address. Alternatively, you can specify the network mask information by entering a forward slash followed by the number of bits in the network mask. For example, you can enter 192.0.0.0 255.255.255.0 as 192.0.0.0/.24.

Configuring the Default Route

You can also assign the default router as the destination by entering 0.0.0.0 0.0.0.0.

Configuring a "Null" Route

You can configure the routing switch to drop IP packets to a specific network or host address by configuring a "null" (sometimes called "null0") static route for the address. When the routing switch receives a packet destined for the address, the routing switch drops the packet instead of forwarding it.

To configure a null static route to drop packets destined for network 209.157.22.x, enter the following commands:

```
HPswitch(config)# ip route 209.157.22.0 255.255.255.0
reject
HPswitch(config)# write memory
```

Syntax: ip route <ip-addr> <ip-mask> reject

or

ip route <ip-addr>/<mask-bits> reject

The <ip-addr> parameter specifies the network or host address. The routing switch will drop packets that contain this address in the destination field instead of forwarding them.

The <ip-mask> parameter specifies the network mask. Ones are significant bits and zeros allow any value. For example, the mask 255.255.255.0 matches on all hosts within the Class C sub-net address specified by <ip-addr>. Alternatively, you can specify the number of bits in the network mask. For example, you can enter 209.157.22.0/24 instead of 209.157.22.0 255.255.255.0.

The **reject** parameter indicates that this is a null route. You must specify this parameter to make this a null route.

Configuring RIP

This section describes how to configure RIP on HP Series 5300XL Switches using the CLI interface.

To display RIP configuration information and statistics, see "Displaying RIP Information" on page 16-26.

Overview of RIP

Routing Information Protocol (RIP) is an IP route exchange protocol that uses a ***distance vector*** (a number representing distance) to measure the cost of a given route. The cost is a distance vector because the cost often is equivalent to the number of router hops between the HP routing switch and the destination network.

An HP routing switch can receive multiple paths to a destination. The software evaluates the paths, selects the best path, and saves the path in the IP route table as the route to the destination. Typically, the best path is the path with the fewest hops. A hop is another router through which packets must travel to reach the destination. If the HP routing switch receives a RIP update from another router that contains a path with fewer hops than the path stored in the HP routing switch's route table, the routing switch replaces the older route with the newer one. The routing switch then includes the new path in the updates it sends to other RIP routers, including HP routing switches.

RIP routers, including HP routing switches, also can modify a route's cost, generally by adding to it, to bias the selection of a route for a given destination. In this case, the actual number of router hops may be the same, but the route has an administratively higher cost and is thus less likely to be used than other, lower-cost routes. A RIP route can have a maximum cost of 15. Any destination with a higher cost is considered unreachable. Although limiting to larger networks, the low maximum hop count prevents endless loops in the network.

HP Series 5300XL Switches support the following RIP types:

- Version 1
- V1 compatible with V2
- Version 2 (the default)

Note

ICMP Host Unreachable Message for Undeliverable ARPs. If the routing switch receives an ARP request packet that it is unable to deliver to the final destination because of the ARP timeout and no ARP response is received (the routing switch knows of no route to the destination address), the routing switch sends an ICMP Host Unreachable message to the source.

RIP Parameters and Defaults

The following tables list the RIP parameters, their default values, and where to find configuration information.

RIP Global Parameters

Table 16-3 lists the global RIP parameters and their default values.

Table 16-3. RIP Global Parameters

Parameter	Description	Default
RIP state	Routing Information Protocol V2-only.	Disabled
auto-summary	Enable/Disable advertisement of summarized routes.	Enabled
metric	Default metric for imported routes.	1
redistribution	RIP can redistribute static and connected routes. (RIP redistributes connected routes by default, when RIP is enabled.)	Disabled

RIP Interface Parameters

Table 16-4 lists the VLAN interface RIP parameters and their default values.

Note

RIP interface configuration is performed on VLAN interfaces on the Series 5300XL Switches.

Table 16-4. RIP Interface Parameters

Parameter	Description	Default
RIP version	The version of the protocol that is supported on the interface. The version can be one of the following: • Version 1 only • Version 2 only • Version 1 compatible with version 2	V2-only
metric	A numeric cost the routing switch adds to RIP routes learned on the interface. This parameter applies only to RIP routes.	1

Parameter	Description	Default
IP address	The routes that a routing switch learns or advertises can be controlled.	The routing switch learns and advertises all RIP routes on all RIP interfaces
loop prevention	The method the routing switch uses to prevent routing loops caused by advertising a route on the same interface as the one on which the routing switch learned the route. • **Split horizon** - the routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route. • **Poison reverse** - the routing switch assigns a cost of 16 ("infinite" or "unreachable") to a route before advertising it on the same interface as the one on which the routing switch learned the route.	Poison reverse
receive	Define the RIP version for incoming packets	V2-only
send	Define the RIP version for outgoing packets	V2-only

Configuring RIP Parameters

Use the following procedures to configure RIP parameters on a system-wide and individual VLAN interface basis.

Enabling RIP

RIP is disabled by default. To enable it, use one of the following methods. When you enable RIP, the default RIP version is **RIPv2-only**. You can change the RIP version on an individual interface basis to **RIPv1** or **RIPv1-compatible-v2** if needed.

To enable RIP on a routing switch, enter the following commands:

```
HPswitch(config)# ip routing
HPswitch(config)# router rip
HPswitch(config-rip-router)# exit
HPswitch(config)# write memory
```

Syntax: [no] router rip

Note

IP routing must be enabled prior to enabling RIP. The first command in the sequence above enables IP routing on the Series 5300XL Switch.

Changing the RIP Type on a VLAN Interface

When you enable RIP on a VLAN interface, **RIPv2-only** is enabled by default. You can change the RIP type to one of the following on an individual VLAN interface basis:

- Version 1 only
- Version 2 only (the default)
- Version 1 - compatible - version 2

To change the RIP type supported on a VLAN interface, enter commands such as the following:

```
HPswitch(config)# vlan 1
HPswitch(vlan-1)# ip rip v1-only
HPswitch(vlan-1)# exit
HPswitch(config)# write memory
```

Syntax: [no] ip rip v1-only | v1-compatible-v2 | v2-only

Changing the Cost of Routes Learned on a VLAN Interface

By default, a Series 5300XL Switch interface increases the cost of a RIP route that is learned on the interface. The switch increases the cost by adding one to the route's metric before storing the route.

You can change the amount that an individual VLAN interface adds to the metric of RIP routes learned on the interface.

Note

RIP considers a route with a metric of 16 to be unreachable. Use this metric only if you do not want the route to be used. In fact, you can prevent the Series 5300XL Switch from using a specific interface for routes learned though that interface by setting its metric to 16.

To increase the cost a VLAN interface adds to RIP routes learned on that interface, enter commands such as the following:

```
HPswitch(config)# interface vlan 1
HPswitch(vlan-1)# ip rip metric 5
```

These commands configure vlan-1 to add 5 to the cost of each route learned on the interface.

Syntax: ip rip metric <1-16>

Configuring RIP Redistribution

You can configure the routing switch to redistribute connected and static routes into RIP. When you redistribute a route into RIP, the routing switch can use RIP to advertise the route to its RIP neighbors.

To configure redistribution, perform the following tasks:

1. Configure redistribution filters to permit or deny redistribution for a route based on the destination network address or interface. (optional)
2. Enable redistribution

Define RIP Redistribution Filters

Route redistribution imports and translates different protocol routes into a specified protocol type. On HP Series 5300XL Switches, redistribution is supported for static routes and directly connected routes only. Redistribution of any other routing protocol into RIP is not currently supported. When you configure redistribution for RIP, you can specify that static or connected routes are imported into RIP routes. Likewise, OSPF redistribution supports the import of static or connected routes into OSPF routes.

To configure for redistribution, define the redistribution tables with "restrict" redistribution filters. In the CLI, use the **restrict** command for RIP at the RIP router level.

Note

Do not enable redistribution until you have configured the redistribution filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

Example: To configure the Series 5300XL Switch to filter out redistribution of static or connected routes on network 10.0.0.0, enter the following commands:

```
HPswitch(config)# router rip
HPswitch(rip)# restrict 10.0.0.0 255.0.0.0
HPswitch(rip)# write memory
```

Note

The default configuration permits redistribution for all default connected routes only.

Syntax: restrict <ip-addr> <ip-mask> | <ip-addr/<prefix length>

This command prevents any routes with a destination address that is included in the range specified by the address/mask pair from being redistributed by RIP.

Modify Default Metric for Redistribution

The default metric is a global parameter that specifies the cost applied to all RIP routes by default. The default value is 1. You can assign a cost from 1 – 15.

Example: To assign a default metric of 4 to all routes imported into RIP, enter the following commands:

```
HPswitch(config)# router rip
HPswitch(rip)# default-metric 4
```

Syntax: default-metric <value>

The **<value>** can be from 1 – 15. The default is 1.

Enable RIP Route Redistribution

Note Do not enable redistribution until you have configured the redistribution filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

To enable redistribution of connected and static IP routes into RIP, enter the following commands.

```
HPswitch(config)# router rip
HPswitch(rip)# redistribute connected
HPswitch(rip)# redistribute static
HPswitch(rip)# write memory
```

Syntax: [no] redistribute connected | static

Changing the Route Loop Prevention Method

RIP can use the following methods to prevent routing loops:

- **Split horizon** - the routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route.
- **Poison reverse** - the routing switch assigns a cost of 16 (“infinity” or “unreachable”) to a route before advertising it on the same interface as the one on which the routing switch learned the route. This is the default.

These loop prevention methods are configurable on an individual VLAN interface basis.

Note These methods are in addition to RIP's maximum valid route cost of 15.

Poison reverse is enabled by default. Disabling poison reverse causes the routing switch to revert to **Split horizon**. (Poison reverse is an extension of Split horizon.) To disable Poison reverse on an interface, and thereby enable Split horizon, enter the following:

```
HPswitch(config)# vlan 1
HPswitch(vlan-1)# no ip rip poison-reverse
```

Syntax: [no] ip rip poison-reverse

Entering the command without the “no” option will re-enable Poison reverse.

Displaying RIP Information

All RIP configuration and status information is shown by the CLI command **show ip rip** and options off that command. The following RIP information can be displayed:

RIP Information Type	Page
General Information	16-27
Interface Information	16-29
Peer Information	16-30
Redistribute Information	16-32
Restrict Information	16-32

Displaying General RIP Information

To display general RIP information, enter the following CLI command at any context level:

```
HPswitch# show ip rip
```

The resulting display will appear similar to the following:

```
RIP global parameters

 RIP protocol    : disabled
 Auto-summary    : disabled
 Default Metric : 1
 Route changes  : 0
 Queries         : 0

RIP interface information

 IP Address      Status      Send mode        Recv mode  Metric      Auth
 --------------- ----------- ---------------- ---------- ----------- ----
 100.1.0.1       enabled     V2-only          V2-only    1           none
 100.2.0.1       enabled     V2-only          V2-only    1           none
 100.3.0.1       enabled     V2-only          V2-only    1           none
 100.4.0.1       enabled     V2-only          V2-only    1           none
 100.10.0.1      enabled     V2-only          V2-only    1           none
 100.11.0.1      enabled     V2-only          V2-only    1           none
 100.12.0.1      enabled     V2-only          V2-only    1           none

RIP peer information

 IP Address      Bad routes  Last update timeticks
 --------------- ----------- ---------------------
```

The display is a summary of Global RIP information, information about interfaces with RIP enabled, and information about RIP peers. The following fields are displayed:

- **RIP protocol** – Status of the RIP protocol on the router. RIP must be enabled here and on the VLAN interface for RIP to be active. The default is **disabled**.
- **Auto-summary** – Status of Auto-summary for all interfaces running RIP. If auto-summary is enabled, then subnets will be summarized to a class network when advertising outside of the given network.

- **Default Metric** – Sets the default metric for imported routes. This is the metric that will be advertised with the imported route to other RIP peers. A RIP metric is a measurement used to determine the 'best' path to network; 1 is the best, 15 is the worse, 16 is unreachable.
- **Route changes** – The number of times RIP has modified the routing switch's routing table.
- **Queries** – The number of RIP queries that have been received by the routing switch.
- **RIP Interface Information** – RIP information on the VLAN interfaces on which RIP is enabled.
  - **IP Address** – IP address of the VLAN interface running rip.
  - **Status** – Status of RIP on the VLAN interface.
  - **Send mode** – The format of the RIP updates: RIP 1, RIP 2, or RIP 2 version 1 compatible.
  - **Recv mode** – The Series 5300XL Switches can process RIP 1, RIP 2, or RIP 2 version 1 compatible update messages.
  - **Metric** – The path "cost", a measurement used to determine the 'best' RIP route path; 1 is the best, 15 is the worse, 16 is unreachable.
  - **Auth** – RIP messages can be required to include an authentication key if enabled on the interface.
- **RIP Peer Information** – RIP Peers are neighboring routers from which the routing switch has received RIP updates.
  - **IP Address** – IP address of the RIP neighbor.
  - **Bad routes** – The number of route entries which were not processed for any reason.
  - **Last update timeticks** – How many seconds have passed since we received an update from this neighbor.

Syntax: `show ip rip`

Displaying RIP Interface Information

To display RIP interface information, enter the following CLI command at any context level:

```
HPswitch# show ip rip interface
```

The resulting display will appear similar to the following:

```
RIP interface information

 IP Address      Status      Send mode        Recv mode  Metric      Auth
 --------------- ----------- ---------------- ---------- ----------- ----
 100.1.0.1       enabled     V2-only          V2-only    1           none
 100.2.0.1       enabled     V2-only          V2-only    1           none
 100.3.0.1       enabled     V2-only          V2-only    1           none
 100.4.0.1       enabled     V2-only          V2-only    1           none
```

See “RIP Interface Information” on the previous page for definitions of these fields.

You can also display the information for a single RIP VLAN interface, by specifying the VLAN ID for the interface, or specifying the IP address for the interface.

Displaying RIP interface information by VLAN ID: For example, to show the RIP interface information for VLAN 1000, enter the following command:

```
HPswitch# show ip rip interface vlan 1000
```

```
RIP interface information for VLAN 1000

 IP Address : 120.1.1.1
 Status     : enabled

 Send mode  : V2-only
 Recv mode  : V2-only
 Metric     : 1
 Auth       : none

 Bad packets received : 0
 Bad routes received  : 0
 Sent updates : 0
```

The information in this display includes the following fields, which are defined under ““RIP Interface Information” on page 16-28: **IP Address**, **Status**, **Send mode**, **Recv mode**, **Metric**, and **Auth**.

The information also includes the following fields:

- **Bad packets received** – The number of packets that were received on this interface and were not processed for any reason.
- **Bad routes received** – The number of route entries that were received on this interface and were not processed for any reason.
- **Sent updates** – The number of RIP routing updates that have been sent on this interface.

Displaying RIP interface information by IP Address: For example, to show the RIP interface information for the interface with IP address 100.2.0.1, enter the following command:

```
HPswitch# show ip rip interface 100.2.0.1
```

```
RIP interface information for 100.2.0.1

  IP Address : 100.2.0.1
  Status     : enabled

  Send mode  : V2-only
  Recv mode  : V2-only
  Metric : 1
  Auth : none

  Bad packets received : 0
  Bad routes received  : 0
  Sent updates : 0
```

The information shown in this display has the same fields as for the display for a specific VLAN ID. See the previous page for the definitions of these fields.

Syntax: show ip rip interface [*ip-addr* | vlan *vlan-id*]

Displaying RIP Peer Information

To display RIP peer information, enter the following CLI command at any context level:

```
HPswitch# show ip rip peer
```

The resulting display will appear similar to the following:

```
RIP peer information

 IP Address      Bad routes  Last update timeticks
 --------------- ----------- ---------------------
 100.1.0.100     0           1
 100.2.0.100     0           0
 100.3.0.100     0           2
 100.10.0.100    0           1
```

This display lists all neighboring routers from which the routing switch has received RIP updates. The following fields are displayed:

- **IP Address** – IP address of the RIP peer neighbor.
- **Bad routes** – The number of route entries that were not processed for any reason.
- **Last update timeticks** – How many seconds have passed since the routing switch received an update from this peer neighbor.

Displaying RIP information for a specific peer: For example, to show the RIP peer information for the peer with IP address 100.1.0.100, enter the following command:

```
HPswitch# show ip rip peer 100.1.0.100
```

```
RIP peer information for 100.0.1.100

 IP Address : 100.1.0.100

 Bad routes  : 0
 Last update timeticks : 2
```

This display lists the following information for a specific RIP peer:

- **IP Address** – IP address of the RIP peer neighbor.
- **Bad routes** – The number of route entries which were not processed for any reason.
- **Last update timeticks** – How many seconds have passed since the routing switch received an update from this neighbor.

Displaying RIP Redistribution Information

To display RIP redistribution information, enter the following CLI command at any context level:

```
HPswitch# show ip rip redistribute
```

```
RIP redistributing

 Route type Status
 ---------- --------
 connected  enabled
 static     enabled
```

RIP automatically redistributes connected routes which are configured on interfaces that are running RIP, and all routes that are learned via RIP. The **router rip redistribute** command, described on page 16-24, configures the routing switch to cause RIP to advertise connected routes that are not running RIP, and static routes. The display shows whether RIP redistribution is enabled or disabled for connected and static routes.

Displaying RIP Redistribution Filter (restrict) Information

To display RIP restrict filter information, enter the following CLI command at any context level:

```
HPswitch# show ip rip restrict
```

```
RIP restrict list

 IP Address      Mask
 --------------- ---------------
```

The display shows if any routes, identified by the IP Address and Mask fields are being restricted from redistribution. The restrict filters are configured by the **router rip restrict** command described on page 16-24.

Configuring OSPF

This section describes how to configure OSPF on HP Series 5300XL Switches using the CLI interface.

To display OSPF configuration information and statistics, see "Displaying OSPF Information" on page 16-52.

Overview of OSPF

OSPF is a link-state routing protocol. The protocol uses link-state advertisements (LSA) to update neighboring routers regarding its interfaces and information on those interfaces. The Series 5300XL Switch floods these LSAs to all neighboring routers to update them regarding the interfaces. Each router maintains an identical database that describes its area topology to help a router determine the shortest path between it and any neighboring router.

HP Series 5300XL Switches support the following types of LSAs, which are described in RFC 2328:

- Router link
- Network link
- Summary link
- Autonomous system (AS) summary link
- AS external link

OSPF is built upon a hierarchy of network components. The highest level of the hierarchy is the ***Autonomous System (AS)***. An autonomous system is defined as a number of networks, all of which share the same routing and administration characteristics.

An AS can be divided into multiple ***areas***. Each area represents a collection of contiguous networks and hosts. Areas define the limit to which link-state advertisements are broadcast, thereby limiting the amount of flooding that occurs within the network. An area is represented in OSPF by either an IP address or a number.

You can further limit the broadcast area of flooding by defining an area range. The area range allows you to assign an aggregate value to a range of IP addresses. This aggregate value becomes the address that is advertised instead all of the individual addresses it represents being advertised. You can assign up to 8 ranges in an OSPF area.

An OSPF router can be a member of multiple areas. Routers with membership in multiple areas are known as ***Area Border Routers (ABRs)***. Each ABR maintains a separate topological database for each area the router is in. Each topological database contains all of the LSA databases for each router within a given area. The routers within the same area have identical topological databases. The ABR is responsible for forwarding routing information or changes between its border areas.

An ***Autonomous System Boundary Router (ASBR)*** is a router that is running multiple protocols and serves as a gateway to routers outside an area and those operating with different protocols. The ASBR is able to import and translate different protocol routes into OSPF through a process known as ***redistribution***. For more details on redistribution and configuration examples, see “Enable Route Redistribution” on page 16-49.

Designated Routers in Multi-Access Networks

In a network that has multiple routers attached, OSPF elects one router to serve as the designated router (DR) and another router on the segment to act as the backup designated router (BDR). This arrangement minimizes the amount of repetitive information that is forwarded on the network by forwarding all messages to the designated router and backup designated routers responsible for forwarding the updates throughout the network.

Designated Router Election

In a network with no designated router and no backup designated router, the neighboring router with the highest priority is elected as the DR, and the router with the next largest priority is elected as the BDR.

If the DR goes off-line, the BDR automatically becomes the DR. The router with the next highest priority becomes the new BDR.

Note

Priority is a configurable option at the interface level. You can use this parameter to help bias one Series 5300XL Switch as the DR.

If two neighbors share the same priority, the router with the highest router ID is designated as the DR. The router with the next highest router ID is designated as the BDR.

Note

By default, the HP router ID is the lowest numbered IP address configured on the device. For more information or to change the router ID, see “Changing the Router ID” on page 16-9.

When multiple Series 5300XL Switches on the same network are declaring themselves as DRs, then both priority and router ID are used to select the designated router and backup designated routers.

When only one router on the network claims the DR role despite neighboring routers with higher priorities or router IDs, this router remains the DR. This is also true for BDRs.

The DR and BDR election process is performed when one of the following events occurs:

- an interface is in a waiting state and the wait time expires
- an interface is in a waiting state and a hello packet is received that addresses the BDR
- a change in the neighbor state occurs, such as:
  - a neighbor state transitions from 2 or higher
  - communication to a neighbor is lost
  - a neighbor declares itself to be the DR or BDR for the first time

OSPF RFC 1583 and 2328 Compliance

HP Series 5300XL Switches are configured, by default, to be compliant with the RFC 1583 OSPF V2 specification. HP Series 5300XL Switches can also be configured to operate with the latest OSPF standard, RFC 2328.

Note

For details on how to configure the system to operate with the RFC 2328, see “Configuring OSPF” on page 16-37.

Reduction of Equivalent AS External LSAs

An OSPF ASBR uses AS External link advertisements (AS External LSAs) to originate advertisements of a route to another routing domain, such as a RIP domain. The ASBR advertises the route to the external domain by flooding AS External LSAs to all the other OSPF routers (except those inside stub networks) within the local OSPF Autonomous System (AS).

In some cases, multiple ASBRs in an AS can originate equivalent LSAs. The LSAs are equivalent when they have the same cost, the same next hop, and the same destination. The HP switch optimizes OSPF by eliminating duplicate AS External LSAs in this case. The switch with the lower router ID flushes the duplicate External LSAs from its database and thus does not flood the duplicate External LSAs into the OSPF AS. AS External LSA reduction therefore reduces the size of the switch’s link state database.

This enhancement implements the portion of RFC 2328 that describes AS External LSA reduction. This enhancement is enabled by default, requires no configuration, and cannot be disabled.

OSPF eliminates duplicate AS External LSAs. When two or more HP Series 5300XL Switches configured as ASBRs have equal-cost routes to the same next-hop router in an external routing domain, the ASBR with the highest router ID floods the AS External LSAs for the external domain into the OSPF AS, while the other ASBRs flush the equivalent AS External LSAs from their databases. As a result, the overall volume of route advertisement traffic within the AS is reduced and the Series 5300 Switches that flush the duplicate AS External LSAs have more memory for other OSPF data.

Algorithm for AS External LSA Reduction. The AS External LSA reduction feature behavior changes under the following conditions:

- There is one ASBR advertising (originating) a route to the external destination, but one of the following happens:
  - A second ASBR comes on-line
  - A second ASBR that is already on-line begins advertising an equivalent route to the same destination.

  In either case above, the Series 5300XL Switch with the higher router ID floods the AS External LSAs and the other Series 5300XL Switch flushes its equivalent AS External LSAs.
- One of the ASBRs starts advertising a route that is no longer equivalent to the route the other ASBR is advertising. In this case, the ASBRs each flood AS External LSAs. Since the LSAs either no longer have the same cost or no longer have the same next-hop router, the LSAs are no longer equivalent, and the LSA reduction feature no longer applies.
- The ASBR with the higher router ID becomes unavailable or is reconfigured so that it is no longer an ASBR. In this case, the other ASBR floods the AS External LSAs.

Dynamic OSPF Activation and Configuration

OSPF is automatically activated when you enable it. The protocol does not require a software reload.

Without ever having to reset the switch, you can change and save all the OSPF configuration options, including the following:

- all OSPF interface-related parameters (for example: area, hello timer, router dead time cost, priority, re-transmission time, transit delay)
- all area parameters
- all area range parameters
- all virtual-link parameters
- all global parameters
- creation and deletion of an area, interface or virtual link
- changes to address ranges
- changes to global values for redistribution
- addition of new virtual links

The only configuration change that requires you to disable and then re-enable OSPF operation is reconfiguring the Router ID.

Configuring OSPF

To begin using OSPF on the Series 5300XL Switch, perform the steps outlined below:

1. Enable routing on the routing switch.
1. Enable OSPF on the routing switch.
2. Assign the areas to which the routing switch will be attached.
3. Assign individual VLAN interfaces to the OSPF areas.
4. Define redistribution “restrict” filters, if desired.
5. Enable redistribution, if you defined redistribution filters.
6. Modify default global and interface parameters as required.
7. Modify OSPF standard compliance, if desired.

Note OSPF is automatically enabled without a system reset.

Configuration Rules

- If a Series 5300 Switch is to operate as an ASBR, you must enable redistribution. When you do that, ASBR capability is automatically enabled.
- All Series 5300 Switch VLAN interfaces on which you wish to run OSPF must be assigned to one of the defined areas. When a VLAN interface is assigned to an area, the primary IP address is automatically included in the assignment. To include secondary addresses, you must enable OSPF on them separately, or use the “all” option in the assignment.

OSPF Parameters

You can modify or set the following global and interface OSPF parameters.

Global Parameters:

- Modify OSPF standard compliance setting.
- Assign an area.
- Define an area range.
- Define the area virtual link.
- Set global default metric for OSPF.
- Define redistribution metric type.
- Enable redistribution.
- Define redistribution restrict filters.
- Modify OSPF Traps generated.

Interface Parameters:

- Assign interfaces to an area.
- Define the authentication key for the interface.
- Modify the cost for a link.
- Modify the dead interval.
- Modify the priority of the interface.
- Modify the retransmit interval for the interface.
- Modify the transit delay of the interface.

Note

When using the CLI, you set global level parameters at the OSPF CONFIG Level of the CLI. To reach that level, enter **router ospf..**at the global CONFIG Level. Interface parameters for OSPF are set at the VLAN CONFIG Level using the CLI command, **ip ospf...**

Enable OSPF on the Series 5300XL Switch

When you enable OSPF on the Series 5300XL Switch, the protocol is automatically activated. To enable OSPF on the Series 5300XL Switch, use the CLI commands:

```
HPswitch(config)# ip routing
HPswitch(config)#router ospf
```

The first command enables routing on the Series 5300XL Switch. The second command launches you into the OSPF router level where you can assign areas and modify OSPF global parameters.

Note

Regarding Disabling OSPF. If you disable OSPF, the Series 5300XL Switch retains all the configuration information for the disabled protocol in flash memory. If you subsequently restart OSPF, that previous configuration will be applied.

Assign OSPF Areas

Once OSPF is enabled on the system, you can assign areas. Assign an IP address or number as the ***area ID*** for each area. The area ID is representative of only the primary IP address. To include secondary addresses, you must enable OSPF on them separately, or use the “all” option in the assignment. Each VLAN interface on a Series 5300XL switch can support 16 areas.

An area can be ***normal*** or a ***stub***.

- **Normal** – Series 5300XL Switches within an OSPF normal area can send and receive External Link State Advertisements (LSAs).
- **Stub** – Series 5300XL Switches within an OSPF stub area cannot send or receive External LSAs. In addition, the routing switches in an OSPF stub area must use a default route to the area’s Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) to send traffic out of the area.

Example: Here is an example of the commands to set up several OSPF areas.

```
HPswitch(ospf)# area 192.5.1.0
HPswitch(ospf)# area 200.5.0.0
HPswitch(ospf)# area 195.5.0.0
HPswitch(ospf)# area 0.0.0.0
HPswitch(ospf)# write memory
```

Syntax: area <num> | <ip-addr> [stub <cost> [no-summary]]

The **<num> | <ip-addr>** parameter specifies the area number, which can be a number or in IP address format. If you specify a number, the number can be from 0 – 4,294,967,295.

The **<cost>** specifies the cost of the default route to be injected into the stub area, if this routing switch is an ABR. The value can be from 1 – 16,777,215. If you configure a stub area, you must specify the cost. There is no default. Normal areas do not use the cost parameter.

The **no-summary** parameter applies only to stub areas and disables summary LSAs from being sent into the area. See "Assign a Totally Stubby Area" below.

Note

You can assign subnets individually to areas. The limit on the number of areas is 16.

Assign a Totally Stubby Area

By default, the Series 5300XL Switch sends summary LSAs (LSA type 3) into stub areas. You can further reduce the number of link state advertisements (LSA) sent into a stub area by configuring the Series 5300XL Switch to stop sending summary LSAs (type 3 LSAs) into the area. You can disable the summary LSAs when you are configuring the stub area or later after you have configured the area.

This feature disables origination of summary LSAs, but the Series 5300XL Switch still accepts summary LSAs from OSPF neighbors and floods them to other neighbors. The Series 5300XL Switch can form adjacencies with other routers regardless of whether summarization is enabled or disabled for areas on each Series 5300XL Switch.

When you enter a command to disable the summary LSAs, the change takes effect immediately. If you apply the option to a previously configured area, the Series 5300XL Switch flushes all of the summary LSAs it has generated (as an ABR) from the area.

Note

This feature applies only when the Series 5300XL Switch is configured as an Area Border Router (ABR) for the area. To completely prevent summary LSAs from being sent to the area, disable the summary LSAs on each OSPF router that is an ABR for the area.

To disable summary LSAs for a stub area, enter commands such as the following:

```
HPswitch(config-ospf-router)# area 40 stub 3 no-summary
```

Syntax: area <num> | <ip-addr> [stub <cost> [no-summary]]

The **<num> | <ip-addr>** parameter specifies the area number, which can be a number or in IP address format. If you specify a number, the number can be from 0 – 4,294,967,295.

The **<cost>** specifies the cost of the default route to be injected into the stub area, if this routing switch is an ABR. The value can be from 1 – 16777215. If you configure a stub area, you must specify the cost. There is no default. Normal areas do not use the cost parameter.

The **no-summary** parameter applies only to stub areas and disables summary LSAs from being sent into the area.

Note

You can assign subnets individually to areas. The limit on the number of areas is 16.

Assigning an Area Range (optional)

You can assign a ***range*** for an area, but it is not required. Ranges allow a specific IP address and mask to represent a range of IP addresses within an area, so that only that reference range address is advertised to the network, instead of all the addresses within that range. Each area can have up to 8 range addresses.

Example. To define an area range for sub-nets on 193.45.5.1 and 193.45.6.2, enter the following commands:

```
HPswitch(config)# router ospf
HPswitch(ospf)# area 192.45.5.1 range 193.45.0.0
255.255.0.0
HPswitch(ospf)# area 193.45.6.2 range 193.45.0.0
255.255.0.0
```

Syntax: area <ospf-area-id | backbone> range <ip-addr/mask-length>
[no-advertise]

The **<ospf-area-id>** parameter specifies the area number, which can be in IP address format.

The **<ip-addr>** parameter following **range** specifies the IP address portion of the range. The software compares the address with the significant bits in the mask. All network addresses that match this comparison are summarized in a single route advertised by the Series 5300XL Switch.

The **<mask-length>** parameter specifies the portions of the IP address that a route must contain to be summarized in the summary route. In the example above, all networks that begin with 209.157 are summarized into a single route.

Assigning VLAN to an Area

Once you define OSPF areas, you can assign VLANs to the areas. All Series 5300XL Switch VLANs must be assigned to one of the defined areas on an OSPF router. When a VLAN is assigned to an area, the primary IP address is automatically included in the assignment. To include secondary addresses, you must enable OSPF on them separately, or use the “all” option in the assignment.

Example: To assign VLAN 8 of Switch A to area 192.5.0.0 and then save the changes, enter the following commands:

```
HPSwitch(ospf)# vlan 8
RouterA(vlan-8)# ip ospf area 192.5.0.0
RouterA(vlan-8)# write memory
```

Modify Interface Defaults

OSPF has interface parameters that you can configure. For simplicity, each of these parameters has a default value.   No change to these default values is required except as needed for specific network configurations.

VLAN default values can be modified using the following CLI commands at the **VLAN interface level** of the CLI:

- ip ospf area <ip-addr>
- ip ospf authentication-key <password>
- ip ospf cost <num>
- ip ospf dead-interval <value>
- ip ospf hello-interval <value>
- ip ospf priority <value>
- ip ospf retransmit-interval <value>
- ip ospf transmit-delay <value>

For a complete description of these parameters, see the summary of OSPF port parameters in the next section.

OSPF Interface Parameters

The following parameters apply to OSPF interfaces.

Area: Assigns an interface to a specific area. You can assign either an IP address or number to represent an OSPF Area ID. If you assign a number, it can be any value from 0 – 4,294,967,295.

Authentication-key: OSPF supports two methods of authentication for each VLAN—none and simple password. Only one method of authentication can be active on a subnet at a time. The default authentication value is none, meaning no authentication is performed.

The simple password method of authentication requires you to configure an alphanumeric password on an interface. The simple password setting takes effect immediately. All OSPF packets transmitted on the interface contain this password. Any OSPF packet received on the interface is checked for this password. If the password is not present, then the packet is dropped. The password can be up to eight characters long.

Cost: Indicates the overhead required to send a packet across an interface. You can modify the cost to differentiate between 100 Mbps and 1000 Mbps (1 Gbps) links. The default cost is always 1.

Dead-interval: Indicates the number of seconds that a neighbor router waits for a hello packet from the current Series 5300XL Switch before declaring the Series 5300XL Switch down. The value can be from 1 – 2,147,483,647 seconds. The default is 40 seconds.

Hello-interval: Represents the length of time between the transmission of hello packets. The value can be from 1 – 65535 seconds. The default is 10 seconds.

Priority: Allows you to modify the priority of an OSPF router. The priority is used when selecting the designated router (DR) and backup designated routers (BDRs). The value can be from 0 – 255. The default is 1. If you set the priority to 0, the Series 5300XL Switch does not participate in DR and BDR election.

Retransmit-interval: The time between retransmissions of link-state advertisements (LSAs) to adjacent routers for this interface. The value can be from 0 – 3600 seconds. The default is 5 seconds.

Transit-delay: The time it takes to transmit Link State Update packets on this interface. The value can be from 0 – 3600 seconds. The default is 1 second.

Assign Virtual Links

It is highly recommended that all ABRs (area border routers) have either a direct or indirect link to the OSPF backbone area (0.0.0.0 or 0). If an ABR does not have a physical link to the area backbone, the ABR can configure a ***virtual link*** to another router within the same area, which has a physical connection to the area backbone.

Note

A backbone area can be purely virtual with no physical backbone links. Also note that virtual links can be “daisy chained”. If so, may not have one end physically connected to the backbone.

The path for a virtual link is through an area shared by the neighbor ABR (router with a physical backbone connection), and the ABR requiring a logical connection to the backbone.

Two parameters fields must be defined for all virtual links—transit area ID and neighbor router.

- The ***transit area ID*** represents the shared area of the two ABRs and serves as the connection point between the two routers. This number should match the area ID value.

- The ***neighbor router*** field is the router ID (IP address) of the router that is physically connected to the backbone, when assigned from the router interface requiring a logical connection. When assigning the parameters from the router with the physical connection, the router ID is the IP address of the router requiring a logical connection to the backbone.

Note

By default, the HP router ID is the lowest numbered IP address configured on the device. For more information or to change the router ID, see "Changing the Router ID" on page 16-9.

Note

When you establish an area virtual link, you must configure it on both of the routers (both ends of the virtual link).

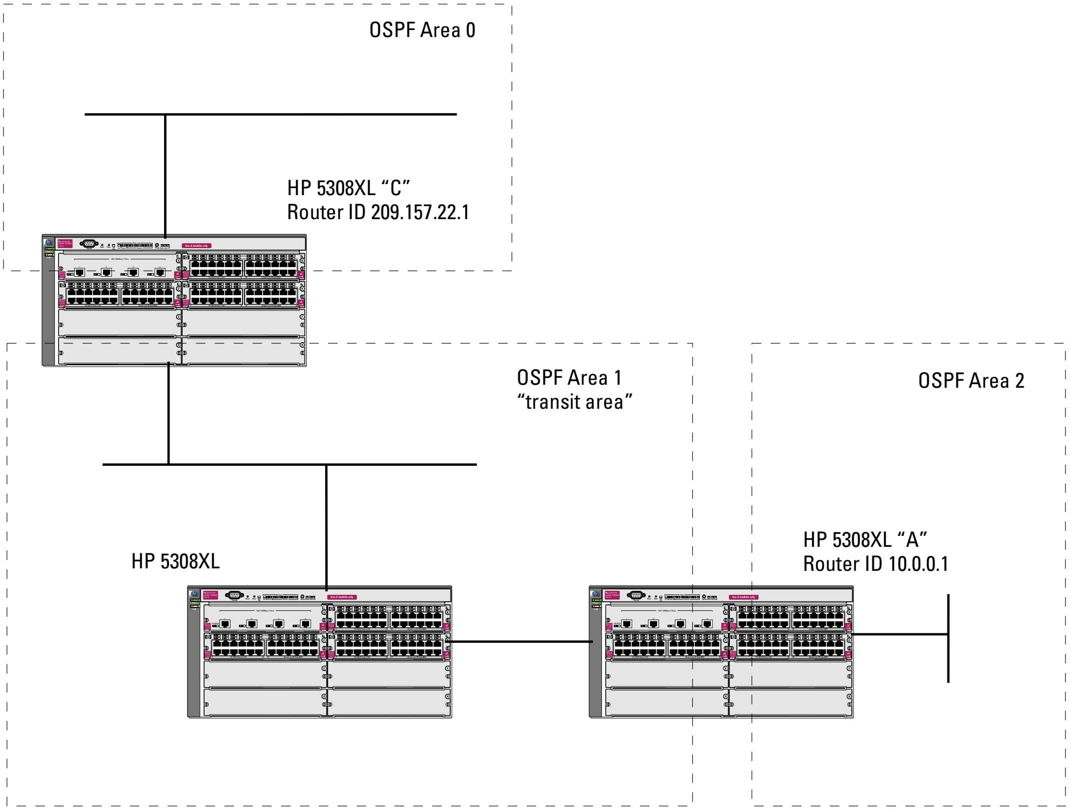

Figure 16-1. Defining OSPF virtual links within a network

Example. Figure 16-1 shows an OSPF area border router, Routing Switch-A, that is cut off from the backbone area (Area 0). To provide backbone access to Routing Switch-A, you can add a virtual link between Routing Switch-A and Routing Switch-C using Area 1 as a transit area. To configure the virtual link, you define the link on the router that is at each end of the link. No configuration for the virtual link is required on the routers in the transit area.

To define the virtual link on Routing Switch-A, enter the following commands:

```
HPswitch(ospf)# area 1 virtual-link 209.157.22.1
HPswitch(ospf)# write memory
```

Enter the following commands to configure the virtual link on Routing Switch-C:

```
HPswitch(ospf)# area 1 virtual-link 10.0.0.1
HPswitch(ospf)# write memory
```

Syntax: area <ip-addr> | <num> virtual-link <router-id> [authentication-key | dead-interval | hello-interval | retransmit-interval | transmit-delay <value>]

The **area <ip-addr> | <num>** parameter specifies the transit area.

The **<router-id>** parameter specifies the router ID of the OSPF router at the remote end of the virtual link. To display the router ID on an HP Series 5300XL Switch, enter the **show ip** command.

See "Modify Virtual Link Parameters" below for descriptions of the optional parameters.

Modify Virtual Link Parameters

OSPF has some parameters that you can modify for virtual links. Notice that these are a subset of the parameters that you can modify for physical interfaces. **cost** is not configured for virtual links, it is calculated by route calculation.

You can modify default values for virtual links using the following CLI command at the ***OSPF router level*** of the CLI, as shown in the following syntax:

Syntax: area <num> | <ip-addr> virtual-link <ip-addr> [authentication-key <string>] [dead-interval <num>] [hello-interval <num>] [retransmit-interval <num>] [transmit-delay <num>]

The parameters are described below. For syntax information, at the CLI prompt, enter the command **area help**.

Virtual Link Parameter Descriptions

You can modify the following virtual link interface parameters:

Authentication Key: This parameter allows you to assign different authentication methods on a per-virtual-link basis. OSPF supports two methods of authentication for each interface—none and simple password. Only one method of authentication can be active on a virtual link at a time.

The simple password method of authentication requires you to configure an alphanumeric password on an interface. The password can be up to eight characters long. The simple password setting takes effect immediately. All OSPF packets transmitted on the interface contain this password. All OSPF packets received on the interface are checked for this password. If the password is not present, then the packet is dropped.

Hello Interval: The length of time between the transmission of hello packets. The range is 1 – 65535 seconds. The default is 10 seconds.

Retransmit Interval: The interval between the re-transmission of link state advertisements to router adjacencies for this interface. The range is 0 – 3600 seconds. The default is 5 seconds.

***Transmit Delay*:** The period of time it takes to transmit Link State Update packets on the interface. The range is 0 – 3600 seconds. The default is 1 second.

Dead Interval: The number of seconds that a neighbor router waits for a hello packet from the current routing switch before declaring the routing switch down. The range is 1 – 65535 seconds. The default is 40 seconds.

Define Redistribution Filters

Route redistribution imports and translates different protocol routes into a specified protocol type. On HP Series 5300XL Switches, redistribution is supported for only static routes and directly connected routes. Redistribution of any other routing protocol into OSPF is not currently supported. When you configure redistribution for OSPF, you can specify that static or connected routes are imported into OSPF routes. Likewise, RIP redistribution supports the import of static or connected routes into RIP routes.

To configure for redistribution, define the redistribution tables with restrict redistribution filters. In the CLI, use the **restrict** command for OSPF at the OSPF router level.

Note

Do not enable redistribution until you have configured the redistribution filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

Example: To configure the Series 5300XL Switch acting as an ASBR to filter out redistribution of static or connected routes on network 10.0.0.0, enter the following commands:

```
HPswitchASBR(config)# router ospf
HPswitch(ospf)# restrict 10.0.0.0 255.0.0.0
HPswitch(ospf)# write memory
```

Note

Redistribution is permitted for all routes by default.

Syntax: restrict <ip-addr> <ip-mask> | <ip-addr/<prefix length>

This command prevents any routes with a destination address that is included in the range specified by the address/mask pair from being redistributed by OSPF.

Modify Default Metric for Redistribution

The default metric is a global parameter that specifies the cost applied to all OSPF routes by default. The default value is 10. You can assign a cost from 1 – 16,777,215.

Example: To assign a default metric of 4 to all routes imported into OSPF, enter the following commands:

```
HPswitch(config)# router ospf
HPswitch(ospf)# default-metric 4
```

Syntax: default-metric <value>

The <value> can be from 1 – 16,777,215. The default is 10.

Enable Route Redistribution

Note

Do not enable redistribution until you have configured the redistribution "restrict" filters. Otherwise, the network might get overloaded with routes that you did not intend to redistribute.

To enable redistribution of connected and static IP routes into OSPF, enter the following commands.

```
HPswitch(config)# router ospf
HPswitch(ospf)# redistribution connected
HPswitch(ospf)# redistribution static
HPswitch(ospf)# write memory
```

Syntax: [no] redistribution connected | static

Modify Redistribution Metric Type

The redistribution metric type is used by default for all routes imported into OSPF. Type 1 metrics are the same "units" as internal OSPF metrics and can be compared directly. Type 2 metrics are not directly comparable, and are treated as larger than the largest internal OSPF metric. The default value is type 2.

To modify the default value to type 1, enter the following command:

```
HPswitch(config-ospf-router)# metric-type type1
```

Syntax: metric-type type1 | type2

The default is **type2**.

Administrative Distance

HP Series 5300XL Switches can learn about networks from various protocols, including RIP, and OSPF. Consequently, the routes to a network may differ depending on the protocol from which the routes were learned. For the Series 5300XL Switches, the administrative distance for OSPF routes is set at 110.

The Series 5300XL Switch selects one route over another based on the source of the route information. To do so, the Series 5300XL Switch can use the administrative distances assigned to the sources.

Modify OSPF Traps Generated

OSPF traps as defined by RFC 1850 are supported on HP Series 5300XL Switches. OSPF trap generation is enabled on the Series 5300XL Switch, by default.

When using the CLI, you can disable all or specific OSPF trap generation by entering the following CLI command:

```
HPswitch(ospf)# no snmp-server trap ospf
```

To later re-enable the trap feature, enter the command:

```
HPswitch(ospf)# snmp-server trap ospf
```

To disable a specific OSPF trap, enter the following command:

```
HPswitch(ospf)# no snmp-server trap ospf <ospf-trap>.
```

These commands are at the OSPF Router Level of the CLI.

Here is a summary of OSPF traps supported on HP Series 5300XL Switches, and their associated MIB objects from RFC 1850:

Table 16-2. OSPF Traps and Associated MIB Objects

OSPF Trap Name	MIB Object
interface-state-change-trap	ospfIfstateChange
virtual-interface-state-change-trap	ospfVirtIfStateChange
neighbor-state-change-trap	ospfNbrStateChange
virtual-neighbor-state-change-trap	ospfVirtNbrStateChange
interface-config-error-trap	ospfIfConfigError
virtual-interface-config-error-trap	ospfVirtIfConfigError
interface-authentication-failure-trap	ospfIfAuthFailure
virtual-interface-authentication-failure-trap	ospfVirtIfAuthFailure
interface-receive-bad-packet-trap	ospfIfrxBadPacket
virtual-interface-receive-bad-packet-trap	ospfVirtIfRxBadPacket
interface-retransmit-packet-trap	ospfTxRetransmit
virtual-interface-retransmit-packet-trap	ospfVirtIfTxRetransmit

OSPF Trap Name	MIB Object
originate-lsa-trap	ospfOriginateLsa
originate-maxage-lsa-trap	ospfMaxAgeLsa
link-state-database-overflow-trap	ospfLsdbOverflow
link-state-database-approaching-overflow-trap	ospfLsdbApproachingOverflow

Examples:

1. To stop an OSPF trap from being collected, use the following CLI command:

```
HPswitch(ospf)#no trap <ospf-trap>
```

2. To disable reporting of the neighbor-state-change-trap, enter the following command:

```
HPswitch(ospf)#no trap neighbor-state-change-trap
```

3. To reinstate the trap, enter the following command:

```
HPswitch(ospf)# trap neighbor-state-change-trap
```

Syntax: [no] snmp-server trap ospf <ospf-trap>

Modify OSPF Standard Compliance Setting

Note

All routes in an AS should be configured with the same compliance setting. If any routers in a domain support only RFC 1583, then all routers must be configured with 1583 compatibility.

If all the routers support RFC 2178 or RFC 2328, you should disable RFC 1583 compatibility on all the routers in the domain, since these standards are more robust against routing loops on external routes.

HP Series 5300XL Switches are configured, by default, to be compliant with the RFC 1583 OSPF V2 specification.

To configure a Series 5300XL Switch to operate with the latest OSPF standard, RFC 2328, enter the following commands:

```
HPswitch(config)# router ospf
HPswitch(ospf)# no rfc1583-compatibility
```

Syntax: [no] rfc1583-compatibility

Displaying OSPF Information

You can use CLI commands to display the following OSPF information:

OSPF Information Type	**Page**
General Information	16-52
Area information	16-54
External link state information	16-55
Interface information	16-56
Link state information	16-59
Neighbor information	16-61
Route information	16-67
Virtual Neighbor information	16-64
Virtual Link information	16-65

Displaying General OSPF Configuration Information

To display general OSPF configuration information, enter the following command at any CLI level:

```
HPswitch# show ip ospf general
```

```
OSPF General Status

  OSPF protocol           : enabled
  Router ID               : 10.0.8.36
  RFC 1583 compatibility  : compatible

  Default import metric      : 1
  Default import metric type : external type 2

  Area Border               : yes
  AS Border                 : yes
  External LSA Count        : 9
  External LSA Checksum Sum : 408218
  Originate New LSA Count   : 24814
  Receive New LSA Count     : 14889
```

Syntax: show ip ospf general

The following fields are shown in the OSPF general status display:

Table 16-3. CLI Display of OSPF General Information

This Field...	Displays...
OSPF protocol	indicates whether OSPF is currently enabled.
Router ID	the Router ID that this routing switch is currently using to identify itself
RFC 1583 compatibility	indicates whether the routing switch is currently using RFC 1583 (compatible) or RFC 2328 (non-compatible rules for calculating external routes.
Default import metric	indicates the default metric that will be used for any routes redistributed into OSPF by this routing switch
Default import metric type	indicates the metric type (type 1 or type 2) that will be used for any routes redistributed into OSPF by this routing switch
Area Border	indicates whether this routing switch is currently acting as an area border router
AS Border	indicates whether this routing switch is currently acting as an autonomous system border router (redistributing routes)
External LSA Count	indicates the total number of external LSAs currently in the routing switch's link state database
External LSA Checksum Sum	the sum of the checksums of all external LSAs currently in the routing switch's link state database (quick check for whether database is in sync with other routers in the routing domain)
Originate New LSA Count	count of the number of times this switch has originated a new LSA
Receive New LSA Count	count of the number of times this switch has received a new LSA

Displaying OSPF Area Information

To display OSPF area information, enter the following command at any CLI level:

```
HPswitch> show ip ospf area

OSPF Area Information

  Area ID         Type   Cost  SPFR   ABR  ASBR LSA   Checksum

  --------------- ------ ----- ------ ---- ---- ----- ----------
  0.0.0.0         normal 0     1      0    0    1     0x0000781f
  192.147.60.0    normal 0     1      0    0    1     0x0000fee6
  192.147.80.0    stub   1     1      0    0    2     0x000181cd
```

Syntax: show ip ospf area [OSPF-AREA-ID]

The [OSPF-AREA-ID] parameter shows information for the specified area. If no area is specified, information for all the OSPF areas configured is displayed.

The OSPF area display shows the following information:

Table 16-4. CLI Display of OSPF Area Information

This Field...	Displays...
Area ID	The identifier for this area.
Type	The area type, which can be one of the following: • normal • stub
Cost	The metric for the default route that the routing switch will inject into a stub area if the routing switch is an ABR for the area. This value only applies to stub areas.
SPFR	The number of times the routing switch has run the shortest path first route calculation for this area.
ABR	The number of area border routers in this area.
ASBR	The number of autonomous system border routers in this area.
LSA	The number of LSAs in the link state database for this area.
Chksum(Hex)	The sum of the checksums of all LSAs currently in the area's link state database. This value can be compared to the value for other routers in the area to verify database synchronization.

Displaying OSPF External Link State Information

To display external link state information, enter the following command at any CLI level:

```
HPswitch# show ip ospf external-link-state
```

When you enter this command, an output similar to the following is displayed:

```
Link State ID   Router ID       Age  Sequence #   Checksum
--------------- --------------- ---- -----------  ----------
10.3.7.0        10.0.8.37       232  0x80000005   0x0000d99f
10.3.8.0        10.0.8.37       232  0x80000005   0x0000cea9
10.3.9.0        10.0.8.37       232  0x80000005   0x0000c3b3
10.3.10.0       10.0.8.37       232  0x80000005   0x0000b8bd
10.3.33.0       10.0.8.36       1098 0x800009cd   0x0000b9dd
```

Syntax: show ip ospf external-link-state

The OSPF external link state display shows the following information:

Table 16-5. CLI Display of OSPF External Link State Information

This Field...	Displays...
Link State ID	LSA ID for this LSA. Normally, the destination of the external route, but may have some “host” bits set.
Router ID	Router ID of the router that originated this external LSA.
Age	Current age (in seconds) of this LSA.
Sequence #	Sequence number of the current instance of this LSA.
Chksum(Hex)	LSA checksum value.

Syntax: show ip ospf external-link-state [status | advertise] [link-state-id <link-state-id> | router-id <router-id> | sequence-number <sequence#>]

The **status** keyword is optional and can be omitted. The output can be filtered to show a subset of the total output by specifying the **link-state-id**, **router-id**, or **sequence-number** options.

The **advertise** keyword displays the hexadecimal data in the specified LSA packet, the actual contents of the LSAs. This can also be filtered as above by including the **link-state-id**, **router-id**, or **sequence-number** options.

An example of the **show ip ospf external-link-state advertise** is the following:

```
OSPF External LSAs

 Advertisements
 ------------------------------------------------------------------------
 000302050a0307000a00082580000005d99f0024ffffff008000000a0000000000000000
 000302050a0308000a00082580000005cea90024ffffff008000000a0000000000000000
 000302050a0309000a00082580000005c3b30024ffffff008000000a0000000000000000
 000302050a030a000a00082580000005b8bd0024ffffff008000000a0000000000000000
 000002050a0321000a000824800009cdb9dd0024ffffff00800000010000000000000000
```

Displaying OSPF Interface Information

To display OSPF interface information, enter the following command at any CLI level:

```
HPswitch# show ip ospf interface
```

```
OSPF Interface Status

 IP Address      Status   Area ID         State   Auth-type Cost   Priority
 --------------- -------- --------------- ------- --------- ------ ---------
 10.3.18.36      enabled  10.3.16.0       BDR     none      1      1
 10.3.53.36      enabled  10.3.48.0       BDR     none      1      1
```

Syntax: show ip ospf interface [vlan <vlan-ID> | <ip-addr>]

The OSPF interface display shows the following information:

Table 16-6. CLI Display of OSPF Interface Information

This Field...	Displays...
IP Address	The local IP address for this interface.
Status	enabled or disabled - whether OSPF is currently enabled on this interface.
Area ID	The ID of the area that this interface is in.

This Field...	Displays...
State	The current state of the interface. The value will be one of the following: • DOWN - the underlying VLAN is down • WAIT - the underlying VLAN is up, but we are waiting to hear hellos from other routers on this interface before we run designated router election • DR - this switch is the designated router for this interface • BDR - this switch is the backup designated router for this interface • DROTHER - this router is not the designated router or backup designated router for this interface
Auth-type	none or simple - will be none if no authentication key is configured, simple if an authentication key is configured. All routers running OSPF on the same link must be using the same authentication type and key.
Cost	The OSPF's metric for this interface.
Priority	This routing switch's priority on this interface for use in the designated router election algorithm.

The **<ip-addr>** parameter displays the OSPF interface information for the specified IP address.

Displaying OSPF Interface Information for a Specific VLAN or IP Address

To display OSPF interface information for a specific VLAN or IP address, enter a command such as the following at any CLI level:

```
HPswitch# show ip ospf interface 10.3.18.36
```

```
OSPF Interface Status for 10.3.18.36

  IP Address      : 10.3.18.36          Status   : enabled
  Area ID         : 10.3.16.0

  State  : BDR                          Auth-type : none
  Cost   : 1                            Priority  : 1
  Type   : BCAST

  Transit Delay     : 1                 Retrans Interval  : 5
  Hello Interval    : 10                Rtr Dead Interval : 40
  Designated Router : 10.3.18.34        Events            : 3
  Backup Desig. Rtr : 10.3.18.36
```

Syntax: show ip ospf interface [vlan <vlan-ID> | <ip-addr>]

The OSPF interface display for a specific VLAN or IP address has the same information as the non-specific show ip ospf interface command for the **IP Address**, **Area ID**, **Status**, **State**, **Auth-type**, **Cost**, and **Priority** fields. See the information for the general command above for definitions of these fields.

The show ip ospf interface command for a specific VLAN or IP address shows the following additional information:

Table 16-7. CLI Display of OSPF Interface Information – VLAN or IP Address

This Field...	Displays...
Type	Will always be **BCAST** for interfaces on this routing switch. Point-to-point or NBMA (frame relay or ATM) type interfaces are not supported on HP Procurve Series 5300XL Switches.
Transit Delay	Configured transit delay for this interface.
Retrans Interval	Configured retransmit interval for this interface.
Hello Interval	Configured hello interval for this interface.
Rtr Dead Interval	Configured router dead interval for this interface.
Designated Router	IP address of the router that has been elected designated router on this interface.

This Field...	Displays...
Backup Desig. Rtr	IP address of the router that has been elected backup designated router on this interface.
Events	Number of times the interface state has changed.

If you issue a **show ip ospf interface vlan <vlan-id>** command, the information will be the same as shown in the previous table, but for the primary IP address on the indicated VLAN.

Displaying OSPF Link State Information

To display OSPF link state information, enter the following command at any CLI level:

```
HPswitch# show ip ospf link-state
```

When you enter this command, an output similar to the following is displayed:

```
OSPF Link State Database for Area 0.0.0.0
                            Advertising
 LSA Type   Link State ID   Router ID       Age  Sequence #  Checksum
 ---------- --------------- --------------- ---- ----------- ----------
 Router     10.0.8.32       10.0.8.32       65   0x80000281  0x0000a7b6
 Router     10.0.8.33       10.0.8.33       1638 0x80000005  0x0000a7c8
 Network    10.3.2.37       10.0.8.37       1695 0x80000006  0x00000443
 Summary    10.3.16.0       10.0.8.33       1638 0x80000007  0x0000c242
 Summary    10.3.16.0       10.0.8.35       1316 0x80000008  0x0000aa58
 Summary    10.3.17.0       10.0.8.33       1638 0x8000027b  0x0000becf
 Summary    10.3.17.0       10.0.8.35       1316 0x80000008  0x0000a957
 AsbSummary 10.0.8.36       10.0.8.33       1412 0x80000002  0x00002cba

OSPF Link State Database for Area 10.3.16.0
                            Advertising
 LSA Type   Link State ID   Router ID       Age  Sequence #  Checksum
 ---------- --------------- --------------- ---- ----------- ----------
 Router     10.0.8.33       10.0.8.33       1727 0x8000027e  0x0000d53c
 Router     10.0.8.34       10.0.8.34       1420 0x80000283  0x0000de4f
 Network    10.3.16.34      10.0.8.34       1735 0x80000005  0x00001465
```

Syntax: show ip ospf link-state

The OSPF link state display shows contents of the LSA database, one table for each area. The following information is shown:

Table 16-8. CLI Display of OSPF Link State Information

This Field...	Displays...
LSA Type	Type of LSA. The possible types are: Router Network Summary AsbSummary
Link State ID	LSA ID for this LSA. The meaning depends on the LSA type.
Advertised Router ID	Router ID of the router that originated this LSA.
Age	Current age (in seconds) of this LSA.
Sequence #	Sequence number of the current instance of this LSA.
Chksum(Hex)	LSA checksum value.

Other options for this command: The **status** keyword is optional and can be omitted. The output can be filtered to show a subset of the total output by specifying the **area-id**, **link-state-id**, **router-id**, LSA **type**, or **sequence-number** options.

The **advertise** keyword displays the hexadecimal data in the specified LSA packet, the actual contents of the LSAs. This can also be filtered as above by including the **area-id**, **link-state-id**, **router-id**, LSA **type**, or **sequence-number** options.

The full syntax of the command is as follows:

Syntax: show ip ospf link-state [status | advertise] [<area-id> | link-state-id <link-state-id> | router-id <router-id> | type <router | network | summary | as-summary> | sequence-number <sequence#>]

An example of the **show ip ospf link-state advertise** is the following:

```
OSPF Link State Database for Area 0.0.0.0

 Advertisements
 ------------------------------------------------------------------------
 000202010a0008200a00082080000281a7b60054000000050a030e00ffffff0003000001...
 000202010a0008210a00082180000006a5c90024010000010a0008230a03112104000002
 000102010a0008230a00082380000015755d006c010000070a030600ffffff0003000001...
 000202020a0302250a0008258000000702440024ffffff000a0008250a0008230a000820
 000202030a0310000a00082180000008c043001cffffff0000000002
 000102030a0310000a00082380000009a859001cffffff0000000001
 000002030a0310000a00082480000009ac53001cffffff0000000002
 000202040a0008240a00082180000032abb001c000000000000000b
 000102040a0008240a00082380000004c12a001c0000000000000002

OSPF Link State Database for Area 10.3.16.0

 Advertisements
 ------------------------------------------------------------------------
 000202010a0008210a0008218000027fd33d0054050000050a031900ffffff0003000001...
 000102010a0008220a0008228000284dc50006000000060a031500ffffff0003000001...
 000102020a0311220a0008228000027bf9080020ffffff000a0008220a000821
```

Displaying OSPF Neighbor Information

To display OSPF neighbor information, enter the following command at any CLI level:

```
HPswitch(ospf)# show ip ospf neighbor
```

```
OSPF Neighbor Information

 Router ID       Pri IP Address      NbIfState State    Rxmt QLen Events
 --------------- --- --------------- --------- -------- --------- ----------
 10.0.8.34       1   10.3.18.34      DR        FULL     0         6
 10.3.53.38      1   10.3.53.38      DR        FULL     0         6
```

Syntax: show ip ospf neighbor [IP-ADDR]

The IP-ADDR can be specified to retrieve detailed information for the specific neighbor only. This is the IP address of the neighbor, not the Router ID.

This display shows the following information.

Table 16-9. CLI Display of OSPF Neighbor Information

Field	Description
Router ID	The router ID of the neighbor.
Pri	The OSPF priority of the neighbor. The priority is used during election of the Designated Router (DR) and Backup designated Router (BDR).
IP Address	The IP address of this routing switch's interface with the neighbor.
NbIfState	The neighbor interface state. The possible values are: • **DR** – this neighbor is the elected designated router for the interface. • **BDR** – this neighbor is the elected backup designated router for the interface. • *blank* – this neighbor is neither the DR or the BDR for the interface.
State	The state of the conversation (the adjacency) between your routing switch and the neighbor. The possible values are: • **INIT** – A Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor. (The Series 5300XL Switch itself did not appear in the neighbor's Hello packet.) All neighbors in this state (or higher) are listed in the Hello packets sent from the associated interface. • **2WAY** – Communication between the two routers is bidirectional. This is the most advanced state before beginning adjacency establishment. The Designated Router and Backup Designated Router are selected from the set of neighbors in the 2Way state or greater. • **EXSTART** – The first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial Database Description (DD) sequence number. Neighbor conversations in this state or greater are called adjacencies. • **EXCHANGE** – The Series 5300XL Switch is describing its entire link state database by sending Database Description packets to the neighbor. Each Database Description packet has a DD sequence number, and is explicitly acknowledged. Only one Database Description packet can be outstanding at any time. In this state, Link State Request packets can also be sent asking for the neighbor's more recent advertisements. All adjacencies in Exchange state or greater are used by the flooding procedure. In fact, these adjacencies are fully capable of transmitting and receiving all types of OSPF routing protocol packets. • **LOADING** – Link State Request packets are sent to the neighbor asking for the more recent advertisements that have been discovered (but not yet received) in the Exchange state. • **FULL** – The neighboring routers are fully adjacent. These adjacencies will now appear in router links and network link advertisements.
`Rxmt QLen`	Remote transmit queue length – the number of LSAs that the routing switch has sent to this neighbor and for which the routing switch is awaiting acknowledgements.
Events	The number of times the neighbor's state has changed.

Displaying OSFPF Redistribution Information

As described under “Enable Route Redistribution” on page 16-49, you can configure the routing switch to redistribute connected and static routes into OSPF. When you redistribute a route into OSPF, the routing switch can use OSPF to advertise the route to its OSPF neighbors.

To display the status of the OSPF redistribution, enter the following command at any CLI context level:

```
HPswitch# show ip ospf redistribute
```

```
OSPF redistributing

 Route type Status
 ---------- --------
 connected  enabled
 static     enabled
```

The display shows whether redistribution of each of the route types, connected and static is enabled.

Displaying OSFPF Redistribution Filter (restrict) Information

As described under “Define Redistribution Filters” on page 16-47, you can configure the redistribution filters on the routing switch to restrict route redistribution by OSPF.

To display the status of the OSPF redistribution filters, enter the following command at any CLI context level:

```
HPswitch# show ip ospf restrict
```

```
OSPF restrict list

 IP Address      Mask
 --------------- ---------------
 10.0.8.0        255.255.248.0
 15.0.0.0        255.0.0.0
```

This display shows the configured restrict entries.

Displaying OSPF Virtual Neighbor Information

To display OSPF virtual neighbor information, enter the following command at any CLI level:

```
HPswitch# show ip ospf virtual-neighbor
```

```
OSPF Virtual Interface Neighbor Information

 Router ID       Area ID         State    IP Address      Events
 --------------- --------------- -------- --------------- --------
 10.0.8.33       10.3.16.0       FULL     10.3.17.33      5
 10.0.8.36       10.3.16.0       FULL     10.3.18.36      5
```

Syntax: show ip ospf virtual-neighbor [area <area-id> | <ip-address>]

This display shows the following information.

Table 16-10. CLI Display of OSPF Virtual Neighbor Information

Field	Description
Router ID	The router ID of this virtual neighbor (configured).
Area ID	The area ID of the transit area for the virtual link to this neighbor (configured).
State	The state of the adjacency with this virtual neighbor. The possible values are the same as the OSPF neighbor states. See the State parameter definition in table 16-9 on page 16-62. Note that virtual neighbors should never stay in the 2WAY state.
IP Address	IP address of the virtual neighbor that the routing switch is using to communicate to that virtual neighbor.
Events	The number of times the virtual neighbor’s state has changed.

Notice from the syntax statement that you can get OSPF virtual neighbor information for a specific area or a specific IP address.

Displaying OSPF Virtual Link Information

To display OSPF virtual link information, enter the following command at any CLI level:

```
HPswitch# show ip ospf virtual-link
```

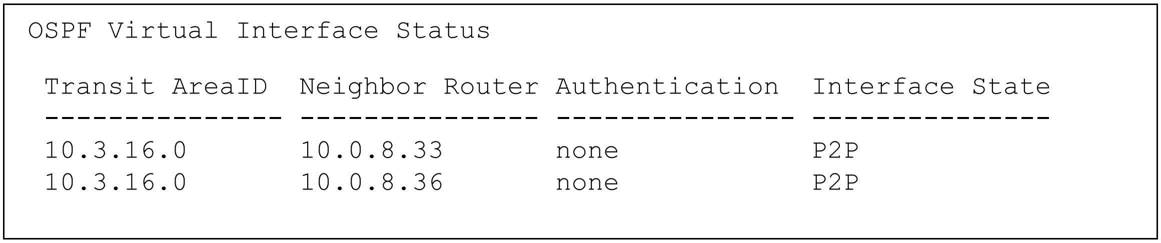

```
OSPF Virtual Interface Status

 Transit AreaID  Neighbor Router Authentication  Interface State
 --------------- --------------- --------------- ---------------
 10.3.16.0       10.0.8.33       none            P2P
 10.3.16.0       10.0.8.36       none            P2P
```

Syntax: show ip ospf virtual-link [area <area-id> | <ip-address>]

This display shows the following information.

Table 16-11. CLI Display of OSPF Virtual Link Information

Field	Description
Transit Area ID	Area ID of transit area for the virtual link.
Neighbor Router	Router ID of the virtual neighbor.
Authentication	none or simple (same as for normal interface).
Interface State	The state of the virtual link to the virtual neighbor. The possible values are: • DOWN – the routing switch has not yet found a route to the virtual neighbor. • P2P – (point-to-point) the routing switch has found a route to the virtual neighbor. Virtual links are "virtual" serial links, hence the point-to-point terminology.

Notice from the syntax statement that you can get OSPF virtual link information for a specific area or a specific IP address.

Example: To get OSPF virtual link information for IP address 10.0.8.33, enter the command:

```
HPswitch# show ip ospf virtual-link 10.0.8.33
```

A display similar to the following is shown.

```
OSPF Virtual Interface Status for interface 10.0.8.33

  Transit AreaID  : 10.3.16.0
  Neighbor Router : 10.0.8.33

  Authentication  : none                 Transit Delay   : 1
  Interface State : P2P                  Rtr Interval    : 5
  Events          : 1                    Hello Interval  : 10
                                         Dead Interval   : 40
```

In this display, these fields show the same type of information as described for the general OSPF virtual link display: **Transit Area ID**, **Neighbor Router**, **Authentication**, and I**nterface State**. This display shows the following additional information:

Table 16-12. CLI Display of OSPF Virtual Link Information – Specific IP Address

Field	**Description**
Events	The number of times the virtual link interface state has changed.
Transit delay	The configured transit delay for the virtual link.
Rtr Interval	The configured retransmit interval for the virtual link.
Hello Interval	The configured hello interval for the virtual link.
Dead Interval	The configured router dead interval for the virtual link

Displaying OSPF Route Information

To display OSPF route and other OSPF configuration information, enter the following command at any CLI level:

```
HPswitch# show ip ospf
```

```
OSPF Configuration Information

  OSPF protocol   : enabled
  Router ID       : 10.0.8.35

 Currently defined areas:

                         Stub          Stub         Stub
  Area ID         Type   Default Cost  Summary LSA  Metric Type
  --------------- ------ ------------- ------------ ---------------
  backbone        normal 1             don't send   ospf metric
  10.3.16.0       normal 1             don't send   ospf metric
  10.3.32.0       normal 1             don't send   ospf metric

 Currently defined address ranges:

  Area ID         LSA Type   IP Network      Network Mask    Advertise
  --------------- ---------- --------------- --------------- ---------
  10.3.16.0       Summary    10.3.16.0       255.255.255.0   yes

 OSPF interface configuration:

                                  Admin          Authen
  IP Address      Area ID         Status   Type  Type   Cost  Pri
  --------------- --------------- -------- ----- ------ ----- ---
  10.3.2.35       backbone        enabled  BCAST none   1     1
  10.3.3.35       backbone        enabled  BCAST none   1     1
  10.3.16.35      10.3.16.0       enabled  BCAST none   1     1
  10.3.32.35      10.3.32.0       enabled  BCAST none   1     1

 OSPF configured interface timers:

                 Transit Retransmit Hello     Dead
  IP Address      Delay   Interval   Interval  Interval
  --------------- ------- ---------- --------- ----------
  10.3.2.35       1       5          10        40
  10.3.3.35       1       5          10        40
  10.3.16.35      1       5          10        40
  10.3.32.35      1       5          10        40

 OSPF configured virtual interfaces:

                                  Authen Xmit   Rxmt   Hello  Dead
  Area ID         Router ID       Type   Delay  Intvl  Intvl  Interval
  --------------- --------------- ------ ------ ------ ------ ----------
  10.3.16.0       10.0.8.33       none   1      5      10     40
  10.3.16.0       10.0.8.36       none   1      5      10     40
```

Syntax: show ip ospf

This screen has a lot of information, most of it already covered in other show commands. The following table shows definitions for the fields:

Table 16-13. CLI Display of OSPF Route and Status Information

Field	Description
OSPF protocol	**enabled** or **disabled** – indicates if OSPF is currently enabled.
Router ID	The Router ID that this routing switch is currently using to identify itself.
Currently Defined Areas:	
Area ID	The identifier for this area.
Type	The type of OSPF area (normal or stub).
Stub Default Cost	The metric for any default route we will inject into a stub area if we are an ABR for the area. This value only applies to stub areas.
Stub Summary LSA	**send** or **don't send** – indicates the state of the no-summary option for the stub area. The value indicates if the area is "totally stubby" (no summaries sent from other areas) or just "stub" (summaries sent). Only applies to stub areas, and only takes effect if the routing switch is the ABR for the area.
Stub Metric Type	This value is always **ospf metric**.
Currently defined address ranges:	
Area ID	The area where the address range is configured.
LSA Type	This value is always **Summary**.
IP Network	The address part of the address range specification.
Network Mask	The mask part of the address range specification.
Advertise	Whether we are advertising (**yes**) or suppressing (**no**) this address range.

Note

The remaining interface and virtual link information is the same as for the previously described OSPF show commands.

Configuring IRDP

The ICMP Router Discovery Protocol (IRDP) is used by HP routing switches to advertise the IP addresses of its router interfaces to directly attached hosts. IRDP is enabled by default. You can enable the feature on a global basis or on an individual VLAN interface basis.

When IRDP is enabled, the routing switch periodically sends Router Advertisement messages out the IP interfaces on which the feature is enabled. The messages advertise the routing switch's IP addresses to directly attached hosts who listen for the messages. In addition, hosts can be configured to query the routing switch for the information by sending Router Solicitation messages.

Some types of hosts use the Router Solicitation messages to discover their default gateway. When IRDP is enabled on the HP routing switch, the routing switch responds to the Router Solicitation messages. Some clients interpret this response to mean that the routing switch is the default gateway. If another router is actually the default gateway for these clients, leave IRDP disabled on the HP routing switch.

IRDP uses the following parameters. If you enable IRDP on individual VLAN interfaces, you can configure these parameters on an individual VLAN interface basis.

- **Packet type** - The routing switch can send Router Advertisement messages as IP broadcasts or as IP multicasts addressed to IP multicast group 224.0.0.1. The default packet type is IP broadcast.
- **Hold time** - Each Router Advertisement message contains a hold time value. This value specifies the maximum about of time the host should consider an advertisement to be valid until a newer advertisement arrives. When a new advertisement arrives, the hold time is reset. The hold time is always longer than the maximum advertisement interval. Therefore, if the hold time for an advertisement expires, the host can reasonably conclude that the router interface that sent the advertisement is no longer available. The default hold time is three times the maximum message interval.
- **Maximum message interval and minimum message interval** - when IRDP is enabled, the routing switch sends the Router Advertisement messages every 450-600 seconds by default. The time within this interval that the routing switch selects is random for each message and is not affected by traffic loads or other network factors. The random interval minimizes the probability that a host will receive Router Advertisement

messages from other routers at the same time. The interval on each IRDP-enabled routing switch interface is independent of the interval on other IRDP-enabled interfaces. The default maximum message interval is 600 seconds. The default minimum message interval is 450 seconds.

- **Preference** - If a host receives multiple Router Advertisement messages from different routers, the host selects the router that send the message with the highest preference as the default gateway. The preference can be a number from -4294967296 to 4294967295. The default is 0.

Enabling IRDP Globally

To enable IRDP globally, enter the following command:

```
HP(config)# ip irdp
```

This command enables IRDP on the IP interfaces on all ports. Each port uses the default values for the IRDP parameters.

Enabling IRDP on an Individual VLAN Interface

To enable IRDP on an individual VLAN interface and configure IRDP parameters, enter commands such as the following:

```
HP(config)# vlan 1
HP(vlan-1)# ip irdp maxadvertinterval 400
```

This example shows how to enable IRDP on a specific interface (VLAN 1) and change the maximum advertisement interval for Router Advertisement messages to 400 seconds.

Syntax: [no] ip irdp [broadcast | multicast] [holdtime <seconds>] [maxadvertinterval <seconds>] [minadvertinterval <seconds>] [preference <number>]

- **broadcast | multicast** - This parameter specifies the packet type the routing switch uses to send the Router Advertisement.
  - **broadcast** - The routing switch sends Router Advertisements as IP broadcasts.
  - **multicast** - The routing switch sends Router Advertisements as multicast packets addressed to IP multicast group 224.0.0.1. This is the default.
- **holdtime** <seconds> - This parameter specifies how long a host that receives a Router Advertisement from the routing switch should consider the advertisement to be valid. When a host receives a new Router Advertisement message from the routing switch, the host resets the hold time

for the routing switch to the hold time specified in the new advertisement. If the hold time of an advertisement expires, the host discards the advertisement, concluding that the router interface that sent the advertisement is no longer available. The value must be greater than the value of the maxadvertinterval parameter and cannot be greater than 9000. The default is three times the value of the maxadvertinterval parameter.

- **maxadvertinterval** - This parameter specifies the maximum amount of time the routing switch waits between sending Router Advertisements. You can specify a value from 1 to the current value of the holdtime parameter. The default is 600 seconds.
- **minadvertinterval** - This parameter specifies the minimum amount of time the routing switch can wait between sending Router Advertisements. The default is three-fourths (0.75) the value of the maxadvertinterval parameter. If you change the maxadvertinterval parameter, the software automatically adjusts the minadvertinterval parameter to be three-fourths the new value of the maxadvertinterval parameter. If you want to override the automatically configured value, you can specify an interval from 1 to the current value of the maxadvertinterval parameter.
- **preference** <number> - This parameter specifies the IRDP preference level of this routing switch. If a host receives Router Advertisements from multiple routers, the host selects the router interface that sent the message with the highest preference as the host's default gateway. The valid range is -4294967296 to 4294967295. The default is 0.

Displaying IRDP Information

To display IRDP information, enter the following command from any CLI level:

```
HPswitch# show ip irdp
```

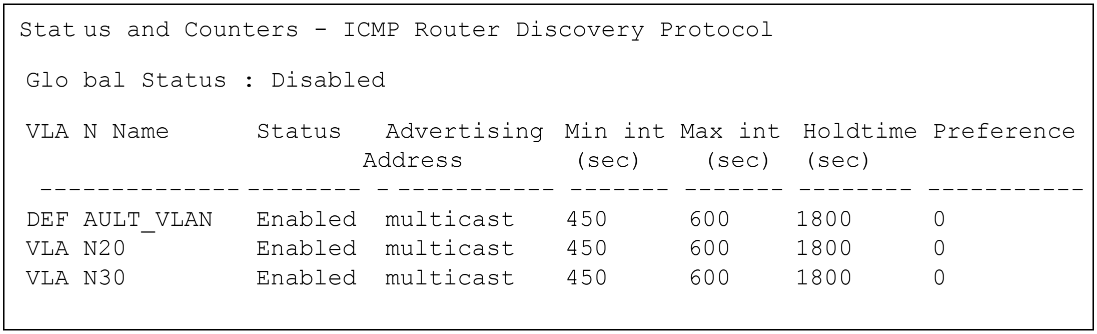

```
Status and Counters - ICMP Router Discovery Protocol

 Global Status : Disabled

 VLAN Name      Status   Advertising  Min int Max int  Holdtime Preference
                         Address       (sec)    (sec)   (sec)
  ------------- -------- - ----------- ------- ------- -------- -----------
 DEFAULT_VLAN   Enabled  multicast    450      600     1800      0
 VLAN20         Enabled  multicast    450      600     1800      0
 VLAN30         Enabled  multicast    450      600     1800      0
```

Configuring DHCP Relay

Overview

The Dynamic Host Configuration Protocol (DHCP) is used for configuring hosts with IP address and other configuration parameters without human intervention. The protocol is composed of three components: the DHCP client, the DHCP server, and the DHCP relay agent. The DHCP client sends broadcast request packets to the network, the DHCP servers respond with broadcast packets that offer IP parameters, such as an IP address for the client. After the client chooses the IP parameters, communication between the client and server is by unicast packets.

The function of the DHCP relay agent is to forward the DHCP messages to other subnets so that the DHCP server doesn't have to be on the same subnet as the DHCP clients. The DHCP relay agent transfers the DHCP messages from DHCP clients located on a subnet without DHCP server, to other subnets. It also relays answers from DHCP servers to DHCP clients.

DHCP Packet Forwarding

The DHCP relay agent on the routing switch forwards DHCP client packets to all DHCP servers that are configured in the table administrated for each VLAN.

Unicast Forwarding

The packets are forwarded using unicast forwarding if the IP address of the DHCP server is a specific host address. The DHCP relay agent sets the destination IP address of the packet to the IP address of the DHCP server and forwards the message.

Broadcast Forwarding

The packets are forwarded using broadcast forwarding if the IP address of the DHCP server is a subnet address or IP broadcast address (255.255.255.255). The DHCP relay agent sets the DHCP server IP address will be set to broadcast IP address and forwarded to all VLANs with configured IP interfaces (except the source VLAN).

Minimum Requirements for DHCP Relay Operation

In order for the DHCP Relay agent to work, the following steps must be completed:

1. DHCP Relay is enabled on the routing switch
2. A DHCP server is servicing the routing switch
3. IP Routing is enabled on the routing switch
4. There is a route from the DHCP server to the routing switch and back
5. An IP Helper address is configured on the routing switch, set to the IP address of the DHCP server on the VLAN that is connected to the DHCP Client.

Enabling DHCP Relay

To enable the DHCP Relay function for the routing switch, at the Config CLI context level, enter the command:

```
HPswitch(config)# dhcp-relay
```

To disable the DHCP Relay function, enter the command:

```
HPswitch(config)# no dhcp-relay
```

Configuring a Helper Address

At the VLAN configuration CLI context level, enter the commands to add the DHCP server's IP address to the VLANs list. For example, to configure a helper address for VLAN 1, enter these commands:

```
HPswitch(conf)# vlan 1
HPswitch(vlan-1)# ip helper-address <ip-addr>
```

To remove the DHCP server helper address, enter this command:

```
HPswitch(vlan-1)# no ip helper-address <ip-addr>
```

A

File Transfers

Contents

Overview

You can download new switch software (operating system—OS) and upload or download switch configuration files. These features are useful for acquiring periodic switch software upgrades and for storing or retrieving a switch configuration.

This appendix includes the following information:

- Downloading an operating system (begins below)
- Transferring switch configurations (begins on page A-13)

For information on how switch memory operates, including primary and secondary flash, see Chapter 5, “Switch Memory and Configuration”.

Downloading an Operating System (OS)

HP periodically provides switch operating system (OS) updates through the HP Procurve website (**http://www.hp.com/go/hpprocurve**). For more information, see the support and warranty booklet shipped with the switch. After you acquire a new OS file, you can use one of the following methods for downloading the operating system (OS) code to the switch:

OS Download Features

Feature	Default	Menu	CLI	Web
TFTP	n/a	page A-4	page A-5	—
Xmodem	n/a	page A-7	page A-8	—
Switch-to-Switch	n/a	page A-9	page A-10	
SNMP Download Manager in HP TopTools for Hubs & Switches	Refer to the documentation provided with HP TopTools for Hubs and Switches			

General OS Download Rules

- An OS image you download via the menu interface always goes to primary flash.
- After an OS download, you must reboot the switch to implement the newly downloaded OS. Until a reboot occurs, the switch continues to run on the OS it was using before the download commenced.

Note

Downloading a new OS does not change the current switch configuration. The switch configuration is contained in separate files that can also be transferred. See “Transferring Switch Configurations” on page A-13.

In most cases, if a power failure or other cause interrupts a flash image download, the switch reboots with the image previously stored in primary flash. In the unlikely event that the primary image is corrupted (which may occur if a download is interrupted by a power failure), the switch goes into boot ROM mode. In this case, use the boot ROM console to download a new image to primary flash. See “Restoring a Flash Image” on page C-35.

Using TFTP To Download an OS Image from a Server

This procedure assumes that:

- An OS file for the switch has been stored on a TFTP server accessible to the switch. (The OS file is typically available from the HP Procurve website at **http://www.hp.com/go/hpprocurve**.)
- The switch is properly connected to your network and has already been configured with a compatible IP address and subnet mask.
- The TFTP server is accessible to the switch via IP.

Before you use the procedure, do the following:

- Obtain the IP address of the TFTP server in which the OS file has been stored.
- If VLANs are configured on the switch, determine the name of the VLAN in which the TFTP server is operating.
- Determine the name of the OS file stored in the TFTP server for the switch (for example, G0101.swi).

Note

If your TFTP server is a Unix workstation, *ensure that the case (upper or lower) that you specify for the filename is the same case as the characters in the OS filenames on the server.*

Menu: TFTP Download from a Server to Primary Flash

Note that the menu interface accesses only the primary flash.

1. In the console Main Menu, select **Download OS** to display this screen:

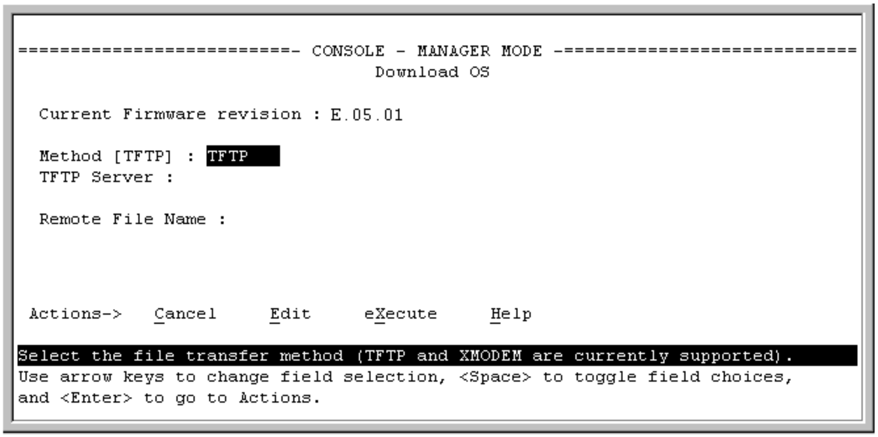

```
=========================- CONSOLE - MANAGER MODE -============================
                              Download OS

  Current Firmware revision : E.05.01

  Method [TFTP]  : TFTP
  TFTP Server :

  Remote File Name :



 Actions->   Cancel     Edit     eXecute     Help

Select the file transfer method (TFTP and XMODEM are currently supported).
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Figure A-1. Example of the Download OS Screen (Default Values)

2. Press [E] (for **Edit**).
3. Ensure that the **Method** field is set to **TFTP** (the default).
4. In the **TFTP Server** field, type in the IP address of the TFTP server in which the OS file has been stored.
5. In the **Remote File Name** field, type the name of the OS file. If you are using a UNIX system, remember that the filename is case-sensitive.
6. Press [Enter], then [X] (for **eXecute**) to begin the OS download. The following screen then appears:

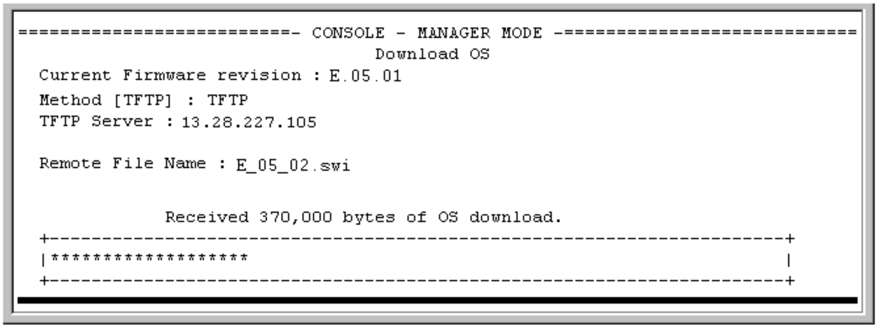

```
=========================- CONSOLE - MANAGER MODE -============================
                              Download OS
  Current Firmware revision : E.05.01
  Method [TFTP]  : TFTP
  TFTP Server : 13.28.227.105

  Remote File Name : E_05_02.swi

              Received 370,000 bytes of OS download.
  +-----------------------------------------------------------------------+
  |*******************                                                    |
  +-----------------------------------------------------------------------+
```

Figure A-2. Example of the Download OS Screen During a Download

A "progress" bar indicates the progress of the download. When the entire OS file has been received, all activity on the switch halts and you will see **Validating and writing system software to FLASH...**

7. After the primary flash memory has been updated with the new operating system, you must reboot the switch to implement the newly downloaded OS. Return to the Main Menu and press [6] (for **Reboot Switch**). You will then see this prompt:

```
Continue reboot of system?  :  No
```

Press the space bar once to change `No` to `Yes`, then press [Enter] to begin the reboot.

Note

When you use the menu interface to download an OS, the new image is always stored in primary flash. Also, using the Reboot Switch command in the Main Menu always reboots the switch from primary flash. Rebooting the switch from the CLI gives you more options. See "Rebooting the Switch" on page 5-17.

8. After you reboot the switch, confirm that the operating system downloaded correctly:
   a. From the Main Menu, select **1. Status and Counters**, and from the Status and Counters menu, select **1. General System Information**
   b. Check the **Firmware revision** line.

CLI: TFTP Download from a Server to Primary or Secondary Flash

This command automatically downloads an OS image to primary or secondary flash.

Syntax: copy tftp flash <*ip-address*> <*remote-os-file*> [< primary | secondary >]

Note that if you do not specify the flash destination, the Xmodem download defaults to primary flash.

For example, to download an OS file named E0500.swi from a TFTP server with the IP address of 10.28.227.103 to primary flash:

1. Execute **copy** as shown below:

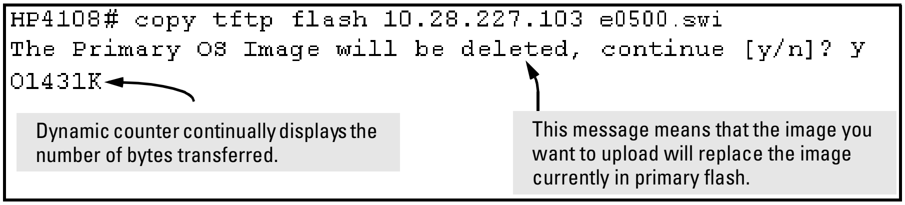

```
HP4108# copy tftp flash 10.28.227.103 e0500.swi
The Primary OS Image will be deleted, continue [y/n]? y
01431K
```

Dynamic counter continually displays the number of bytes transferred.

This message means that the image you want to upload will replace the image currently in primary flash.

Figure A-3. Example of the Command to Download an OS

2. When the switch finishes downloading the OS file from the server, it displays this progress message:

   Validating and Writing System Software to FLASH . . .

3. When the download finishes, you must reboot the switch to implement the newly dowloaded OS. To do so, use one of the following commands:

   boot system flash <primary | secondary> Reboots from the selected flash.

   -or-

   reload Reboots from the flash image currently in use.

   (For more on these commands, see "Rebooting the Switch" on page 5-17.)

4. To confirm that the operating system downloaded correctly, execute **show system** and check the Firmware revision line.

If you need information on primary/secondary flash memory and the boot commands, see "Using Primary and Secondary Flash Image Options" on page 5-12.

Using Xmodem to Download an OS Image From a PC or UNIX Workstation

This procedure assumes that:

- The switch is connected via the Console RS-232 port to a PC operating as a terminal. (Refer to the *Installation and Getting Started Guide* you received with the switch for information on connecting a PC as a terminal and running the switch console interface.)
- The switch operating system (OS) is stored on a disk drive in the PC.

- The terminal emulator you are using includes the Xmodem binary transfer feature. (For example, in the HyperTerminal application included with Windows NT, you would use the **Send File** option in the **Transfer** dropdown menu.)

Menu: Xmodem Download to Primary Flash

Note that the menu interface accesses only the primary flash.

1. From the console Main Menu, select

   7. Download OS

2. Press [E] (for **Edit**).
3. Use the Space bar to select **XMODEM** in the **Method** field.
4. Press [Enter], then [X] (for **eXecute**) to begin the OS download. The following message then appears:

   Press enter and then initiate Xmodem transfer
   from the attached computer.....

5. Press [Enter] and then execute the terminal emulator command(s) to begin Xmodem binary transfer. For example, using HyperTerminal:
   a. Click on **Transfer**, then **Send File**.
   b. Type the file path and name in the Filename field.
   c. In the Protocol field, select **Xmodem**.
   d. Click on the [Send] button.

   The download will then commence. It can take several minutes, depending on the baud rate set in the switch and in your terminal emulator.

6. After the primary flash memory has been updated with the new operating system, you must reboot the switch to implement the newly downloaded OS. Return to the Main Menu and press [6] (for **Reboot Switch**). You will then see this prompt:

   ```
   Continue reboot of system?  :  No
   ```

   Press the space bar once to change `No` to `Yes`, then press [Enter] to begin the reboot.

7. To confirm that the operating system downloaded correctly:
   a. From the Main Menu, select

      1. Status and Counters

         1. General System Information

   b. Check the **Firmware revision** line.

CLI: Xmodem Download from a PC or Unix Workstation to Primary or Secondary Flash

Using Xmodem and a terminal emulator, you can download an OS image to either primary or secondary flash.

Syntax: copy xmodem flash [<primary | secondary>]

Note that if you do not specify the flash destination, the Xmodem download defaults to primary flash.

For example, to download an OS file named E0500.swi from a PC (running a terminal emulator program such as HyperTerminal) to primary flash:

1. Execute the following command in the CLI:

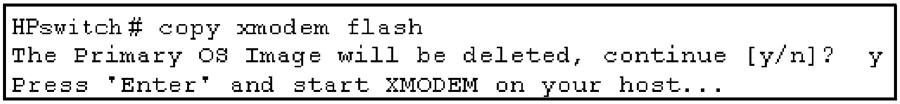

```
HPswitch# copy xmodem flash
The Primary OS Image will be deleted, continue [y/n]?  y
Press 'Enter' and start XMODEM on your host...
```

2. Execute the terminal emulator commands to begin the Xmodem transfer. For example, using HyperTerminal:

   a. Click on **Transfer**, then **Send File**.

   b. Type the file path and name in the Filename field.

   c. In the Protocol field, select **Xmodem**.

   d. Click on the Send button.

   The download can take several minutes, depending on the baud rate used in the transfer.

3. When the download finishes, you must reboot the switch to implement the newly dowloaded OS. To do so, use one of the following commands:

   boot system flash <primary | secondary> Reboots from the selected flash.

   -or-

   reload Reboots from the flash image currently in use.

   (For more on these commands, see “Rebooting the Switch” on page 5-17.)

4. To confirm that the operating system downloaded correctly:

   ```
   HPswitch> show system
   ```

   Check the **Firmware revision** line. It should show the OS version that you downloaded in the preceding steps.

If you need information on primary/secondary flash memory and the boot commands, see “Using Primary and Secondary Flash Image Options” on page 5-12.

Switch-to-Switch Download

You can use TFTP to transfer an OS image between two Series 5300XL switches. The menu interface enables you to transfer primary-to-primary or secondary-to-primary. The CLI enables all combinations of flash location options.

Menu: Switch-to-Switch Download to Primary Flash

Using the menu interface, you can download an OS from either the primary or secondary flash of one Series 5300XL switch to the primary flash of another Series 5300XL switch.

1. From the switch console Main Menu in the switch to receive the download, select **7. Download OS** screen.
2. Ensure that the **Method** parameter is set to **TFTP** (the default).
3. In the **TFTP Server** field, enter the IP address of the remote Series 5300XL switch containing the OS you want to download.
4. For the **Remote File Name**, enter one of the following:
   - To download the OS in the primary flash of the source switch, type “**flash**” in lowercase characters.
   - To download the OS in the secondary flash of the source switch, type **/os/secondary**.
5. Press [Enter], then [X] (for **eXecute**) to begin the OS download.
6. A “progress” bar indicates the progress of the download. When the entire operating system has been received, all activity on the switch halts and the following messages appear:

   Validating and writing system software to FLASH...

7. After the primary flash memory has been updated with the new operating system, you must reboot the switch to implement the newly downloaded OS. Return to the Main Menu and press [6] (for **Reboot Switch**). You will then see this prompt:

   ```
   Continue reboot of system?  :  No
   ```

   Press the space bar once to change `No` to `Yes`, then press [Enter] to begin the reboot.

8. To confirm that the operating system downloaded correctly:
   a. From the Main Menu, select

      Status and Counters

         General System Information

   b. Check the **Firmware revision** line.

CLI: Switch-To-Switch Downloads

You can download an OS image between two Series 5300XL switchs connected on your LAN by initiating a **copy tftp** command from the destination switch.The options for this CLI feature include:

- Copy from primary flash in the source to either primary or secondary in the destination.
- Copy from either primary or secondary flash in the source to either primary or secondary flash in the destination.

Downloading from Primary Only. This command (executed in the destination switch) downloads the OS flash in the source switch's primary flash to either the primary or secondary flash in the destination switch.

Syntax: copy tftp flash <*ip-addr*> flash [primary | secondary]

If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

For example, to download an OS file from primary flash in a Series 5300XL switch with an IP address of 10.28.227.103 to the primary flash in the destination switch, you would execute the following command in the destination switch's CLI:

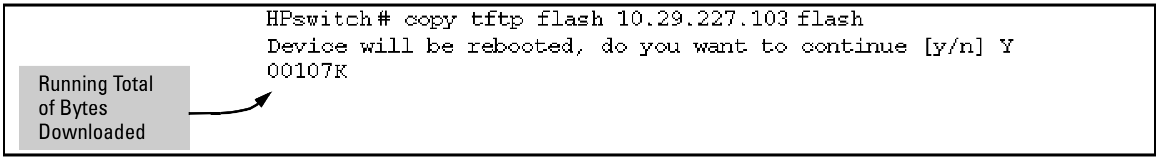

```
HPswitch# copy tftp flash 10.29.227.103 flash
Device will be rebooted, do you want to continue [y/n] Y
00107K
```

Figure A-4. Switch-To-Switch, from Primary in Source to Either Flash in Destination

Downloading from Either Flash in the Source Switch to Either Flash in the Destination Switch. This command (executed in the destination switch) gives you the most options for downloading between switches.

Syntax: copy tftp flash <*ip-addr*> </os/primary> | </os/secondary>
[primary | secondary]

If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

For example, to download an OS file from secondary flash in a Series 5300XL switch with an IP address of 10.28.227.103 to the secondary flash in the destination switch, you would execute the following command in the destination switch's CLI:

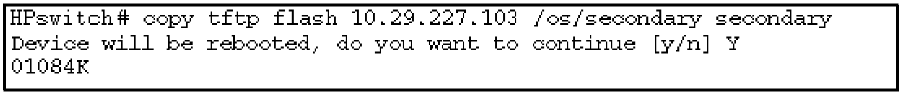

```
HPswitch# copy tftp flash 10.29.227.103 /os/secondary secondary
Device will be rebooted, do you want to continue [y/n] Y
01084K
```

Figure A-5. Switch-to-Switch, from Either Flash in Source to Either Flash in Destination

Using the HP TopTools for Hubs & Switches Utility

HP TopTools for Hubs & Switches includes a software update utility for updating on HP ProCurve switch products such as the Series 2500 switches. For further information, refer to the *HP TopTools for Hubs & Switches User Guide*, provided electronically with the HP TopTools software.

Troubleshooting TFTP Downloads

When using the menu interface, if a TFTP download fails, the Download OS screen indicates the failure.

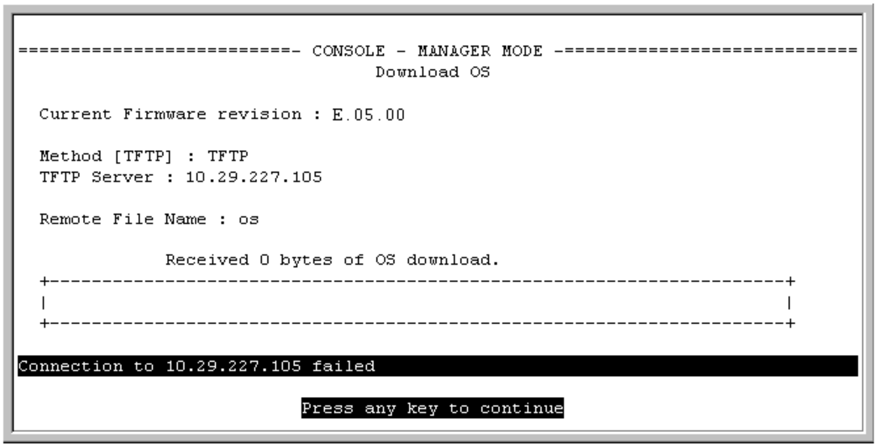

```
==========================- CONSOLE - MANAGER MODE -============================
                                  Download OS

  Current Firmware revision : E.05.00

  Method [TFTP] : TFTP
  TFTP Server : 10.29.227.105

  Remote File Name : os

              Received 0 bytes of OS download.
  +-----------------------------------------------------------------------+
  |                                                                       |
  +-----------------------------------------------------------------------+

Connection to 10.29.227.105 failed

                          Press any key to continue
```

Message Indicating cause of TFTP Download Failure

Figure A-6. Example of Message for Download Failure

To find more information on the cause of a download failure, examine the messages in the switch's Event Log by executing this CLI command:

```
HPswitch# show log tftp
```

(For more on the Event Log, see "Using the Event Log To Identify Problem Sources" on "Using the Event Log To Identify Problem Sources" on page C-22.)

Some of the causes of download failures include:

- Incorrect or unreachable address specified for the **TFTP Server** parameter. This may include network problems.
- Incorrect VLAN.
- Incorrect name specified for the **Remote File Name** parameter, or the specified file cannot be found on the TFTP server. This can also occur if the TFTP server is a Unix machine and the case (upper or lower) for the filename on the server does not match the case for the filename entered for the **Remote File Name** parameter in the Download OS screen.
- One or more of the switch's IP configuration parameters are incorrect.

- For a Unix TFTP server, the file permissions for the OS file do not allow the file to be copied.
- Another console session (through either a direct connection to a terminal device or through Telnet) was already running when you started the session in which the download was attempted.

Note

If an error occurs in which normal switch operation cannot be restored, the switch automatically reboots itself. In this case, an appropriate message is displayed after the switch reboots.

Transferring Switch Configurations

Transfer Features

Feature	Default	Menu	CLI	Web
use TFTP to copy from a remote host to a config file	n/a	—	below	
use TFTP to copy a config file to a remote host	n/a	—	page A-14	
use Xmodem to copy a configuration from a serially connected host to a config file	n/a	—	page A-14	
Use Xmodem to copy a config file to a serially connected host	n/a	—	page A-15	

Using the CLI commands described in this section, you can copy switch configurations to and from a switch.

TFTP: Copying a Configuration from a Remote Host.

Syntax: copy tftp <startup-config | running-config> <*ip-address*> <*remote-file*>

This command copies a configuration from a remote host to the startup-config file in the switch. (See Chapter 5, “Using Primary and Secondary Flash Image Options” for information on the startup-config file.)

For example, to download a configuration file named **sw5300** in the **configs** directory on drive “**d**” in a remote host having an IP address of 10.28.227.105:

```
Series 5300XL switch# copy tftp startup-config
10.28.227.105
          d:\configs\sw2512
```

TFTP: Copying a Configuration File to a Remote Host.

Syntax: copy <startup-config | running-config> tftp <*ip-addr*> <*remote-file*>

This command copies the switch's startup configuration (startup-config file) to a remote TFTP host.

For example, to upload the current startup configuration to a file named **sw5300** in the configs directory on drive "**d**" in a remote host having an IP address of 10.28.227.105:

```
HPswitch# copy startup-config tftp 10.28.227.105
          d:\configs\sw5300
```

Xmodem: Copying a Configuration File from the Switch to a Serially Connected PC or Unix Workstation. To use this method, the switch must be connected via the serial port to a PC or Unix workstation to which you want to copy the configuration file. You will need to:

- Determine a filename to use.
- Know the directory path you will use to store the configuration file.

Syntax: copy <startup-config | running-config> xmodem <pc | unix>

For example, to copy a configuration file to a PC serially connected to the switch:

1. Determine the file name and directory location on the PC.
2. Execute the following command:

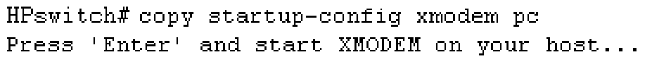

```
HPswitch# copy startup-config xmodem pc
Press 'Enter' and start XMODEM on your host...
```

3. After you see the above prompt, press Enter.
4. Execute the terminal emulator commands to begin the file transfer.

Xmodem: Copying a Configuration File from a Serially Connected PC or Unix Workstation. To use this method, the switch must be connected via the serial port to a PC or Unix workstation on which is stored the configuration file you want to copy. To complete the copying, you will need to know the name of the file to copy and the drive and directory location of the file.

Syntax: copy xmodem startup-config <pc | unix>

For example, to copy a configuration file from a PC serially connected to the switch:

1. Execute the following command:

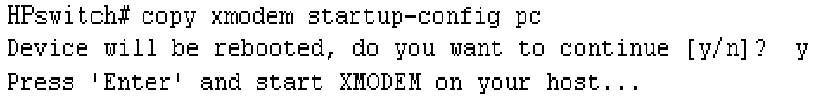

```
HPswitch# copy xmodem startup-config pc
Device will be rebooted, do you want to continue [y/n]?  y
Press 'Enter' and start XMODEM on your host...
```

2. After you see the above prompt, press Enter.
3. Execute the terminal emulator commands to begin the file transfer.
4. When the download finishes, you must reboot the switch to implement the newly dowloaded OS. To do so, use one of the following commands:

boot system flash <primary | secondary>	Reboots from the selected flash.
-or-	
reload	Reboots from the flash image currently in use.

(For more on these commands, see “Rebooting the Switch” on page 5-17.)

Copying Diagnostic Data to a Remote Host, PC, or Unix Workstation

You can use the CLI to copy the following types of switch data to a text file in a management device:

- Command Output: Sends the output of a switch CLI command as a file on the destination device.
- Event Log: Copies the switch's Event Log into a file on the destination device.
- Crash Data: OS-specific data useful for determining the reason for a system crash.
- Crash Log: Processor-Specific operating data useful for determining the reason for a system crash.

Copying Command Output to a Destination Device

This command directs the displayed output of a CLI command to a file in a destination device.

Syntax: copy command-output <"*cli-command*"> tftp <*ip-address*> <*filepath-filename*>

copy command-output <"*cli-command*"> xmodem

For example, to use Xmodem to copy the output of **show config** to a serially connected PC:

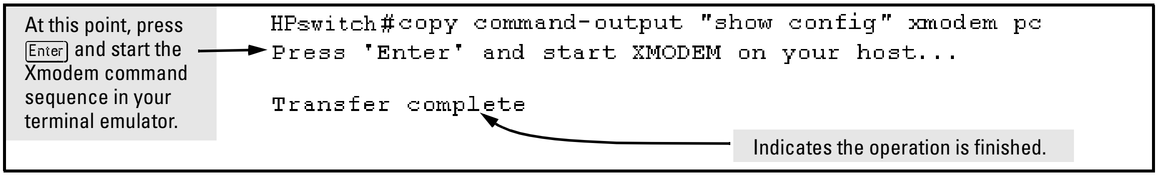

```
HPswitch#copy command-output "show config" xmodem pc
Press 'Enter' and start XMODEM on your host...

Transfer complete
```

At this point, press Enter and start the Xmodem command sequence in your terminal emulator.

Indicates the operation is finished.

Figure A-7. Example of Sending Command Output to a File on an Attached PC

Note that the command you specify must be enclosed in double-quote marks.

Copying Event Log Output to a Destination Device

This command uses TFTP or Xmodem to copy the Event Log content to a PC or UNIX workstation on the network.

Syntax: copy event-log tftp <*ip-address*> <*filepath and filename*>

copy event-log xmodem

For example, to copy the event log to a PC connected to the switch:

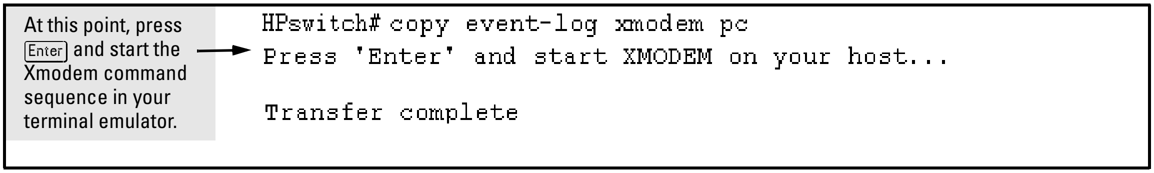

At this point, press [Enter] and start the Xmodem command sequence in your terminal emulator.

```
HPswitch# copy event-log xmodem pc
Press 'Enter' and start XMODEM on your host...

Transfer complete
```

Figure A-8. Example of Sending Event Log Content to a File on an Attached PC

Copying Crash Data Content to a Destination Device

This command uses TFTP or Xmodem to copy the Crash Data content to a PC or UNIX workstation on the network. You can copy individual slot information or the master switch information. If you do not specify either, the command defaults to the master data.

Syntax: copy crash-data [<*slot-id* | master>] xmodem
copy crash-data [<*slot-id* | master>] tftp <*ip-address*> <*filename*>

where: *slot-id* = **a** - **h**, and retrieves the crash log or crash data from the processor on the module in the specified slot.

master Retrieves crash log or crash data from the switch's chassis processor.

For example, to copy the switch's crash data to a file in a PC:

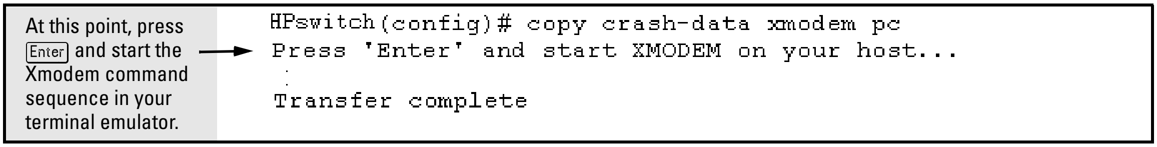

At this point, press [Enter] and start the Xmodem command sequence in your terminal emulator.

```
HPswitch(config)# copy crash-data xmodem pc
Press 'Enter' and start XMODEM on your host...
 :
Transfer complete
```

Figure A-9. Example of Copying Switch Crash Data Content to a PC

Copying Crash Log Data Content to a Destination Device

This command uses TFTP or Xmodem to copy the Crash Log content to a PC or UNIX workstation on the network. You can copy individual slot information or the master switch information. If you do not specify either, the command defaults to the master data.

Syntax: copy crash-log [<*slot-id* | master>] tftp <*ip-address*> <*filepath and filename*>

copy crash-log [<*slot-id* | master>] xmodem

where: *slot-id* = **a** - **h**, and retrieves the crash log or crash data from the processor on the module in the specified slot.

master Retrieves crash log or crash data from the switch's chassis processor.

For example, to copy the Crash Log for slot C to a file in a PC connected to the switch:

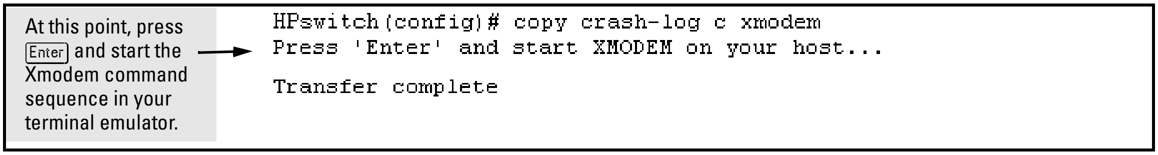

At this point, press Enter and start the Xmodem command sequence in your terminal emulator.

```
HPswitch(config)# copy crash-log c xmodem
Press 'Enter' and start XMODEM on your host...

Transfer complete
```

Figure A-10. Example of sending a Crash Log for Slot C to a File on an Attached PC

Monitoring and Analyzing Switch Operation

Contents

Overview

The Series 5300XL switches have several built-in tools for monitoring, analyzing, and troubleshooting switch and network operation:

- **Status:** Includes options for displaying general switch information, management address data, port status, port and trunk group statistics, MAC addresses detected on each port or VLAN, and STP, IGMP, and VLAN data *(page B-3)*.
- **Counters:** Display details of traffic volume on individual ports *(page B-9)*.
- **Event Log**: Lists switch operating events *("Using the Event Log To Identify Problem Sources" on page C-22)*.
- **Alert Log:** Lists network occurrences detected by the switch—in the Status | Overview screen of the web browser interface *(page 4-15)*.
- **Configurable trap receivers:** Uses SNMP to enable management stations on your network to receive SNMP traps from the switch *(page 10-8)*.
- **Port monitoring (mirroring):** Copy all traffic from the specified ports to a designated monitoring port *(page B-22)*.

Note

Link test and ping test—analysis tools in troubleshooting situations—are described in appendix C, "Troubleshooting". See page C-27.

Status and Counters Data

This section describes the status and counters screens available through the switch console interface and/or the web browser interface.

Note

You can access all console screens from the web browser interface via Telnet to the console. Telnet access to the switch is available in the Device View window under the **Configuration** tab.

Status or Counters Type	Interface	Purpose	Page
Menu Access to Status and Counters	Menu	Access menu interface for status and counter data.	**B-4**
General System Information	Menu, CLI	Lists switch-level operating information.	**B-5**
Management Address Information	Menu, CLI	Lists the MAC address, IP address, and IPX network number for each VLAN or, if no VLANs are configured, for the switch.	**B-6**
Module Information	Menu, CLI	Lists the module type and description for each slot in which a module is installed.	**B-7**
Port Status	Menu, CLI, Web	Displays the operational status of each port.	**B-8**
Port and Trunk Statistics and Flow Control Status	Menu, CLI, Web	Summarizes port activity and lists per-port flow control status.	**B-9**
VLAN Address Table	Menu, CLI	Lists the MAC addresses of nodes the switch has detected on specific VLANs, with the corresponding switch port.	**B-11**
Port Address Table	Menu, CLI	Lists the MAC addresses that the switch has learned from the selected port.	**B-11**
STP Information	Menu, CLI	Lists Spanning Tree Protocol data for the switch and for individual ports. If VLANs are configured, reports on a per-VLAN basis.	**B-16**
IGMP Status	Menu, CLI	Lists IGMP groups, reports, queries, and port on which querier is located.	**B-18**
VLAN Information	Menu, CLI	For each VLAN configured in the switch, lists 802.1Q VLAN ID and up/down status.	**B-19**
Port Status Overview and Port Counters	Web	Shows port utilization and counters, and the Alert Log.	**B-21**

Menu Access To Status and Counters

Beginning at the Main Menu, display the Status and Counters menu by selecting:

1. Status and Counters

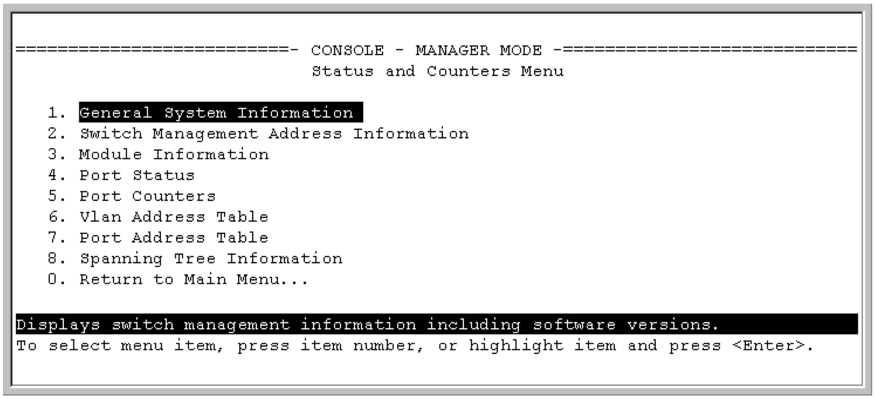

```
==========================- CONSOLE - MANAGER MODE -============================
                           Status and Counters Menu

   1. General System Information
   2. Switch Management Address Information
   3. Module Information
   4. Port Status
   5. Port Counters
   6. Vlan Address Table
   7. Port Address Table
   8. Spanning Tree Information
   0. Return to Main Menu...

Displays switch management information including software versions.
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure B-1. The Status and Counters Menu

Each of the above menu items accesses the read-only screens described on the following pages. Refer to the online help for a description of the entries displayed in these screens.

General System Information

Menu Access

From the console Main Menu, select:

1. Status and Counters

1. General System Information

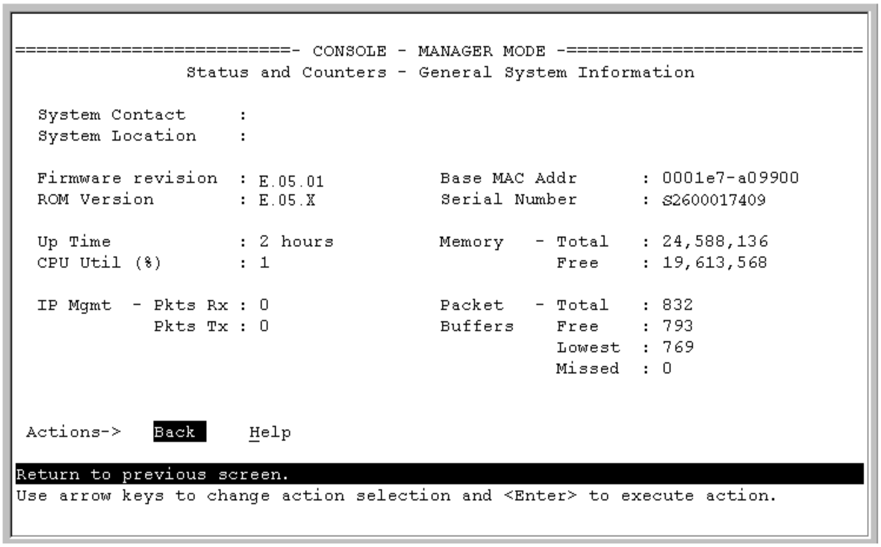

```
=========================- CONSOLE - MANAGER MODE -===========================
                 Status and Counters - General System Information

  System Contact     :
  System Location    :

  Firmware revision  : E.05.01          Base MAC Addr      : 0001e7-a09900
  ROM Version        : E.05.X           Serial Number      : SG2600017409

  Up Time            : 2 hours          Memory   - Total   : 24,588,136
  CPU Util (%)       : 1                           Free    : 19,613,568

  IP Mgmt  - Pkts Rx : 0                Packet   - Total   : 832
             Pkts Tx : 0                Buffers    Free    : 793
                                                   Lowest  : 769
                                                   Missed  : 0


 Actions->   Back     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure B-2. Example of General Switch Information

This screen dynamically indicates how individual switch resources are being used. See the online Help for details.

CLI Access

Syntax: show system-information

Switch Management Address Information

Menu Access

From the Main Menu, select:

1 Status and Counters . . .

2. Switch Management Address Information

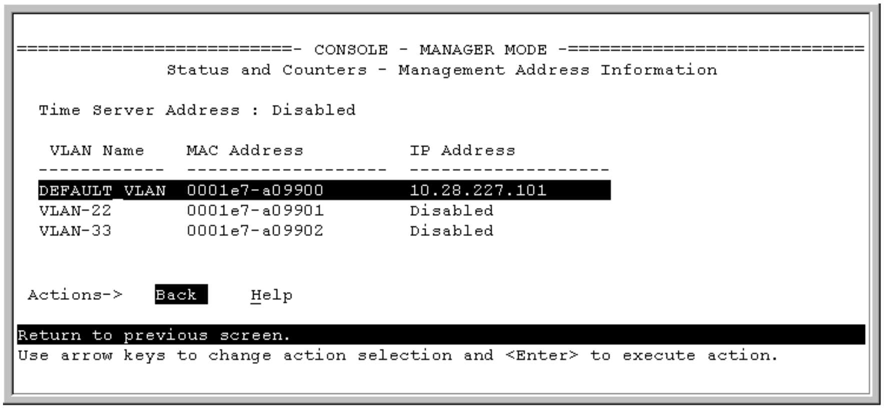

```
========================- CONSOLE - MANAGER MODE -===========================
              Status and Counters - Management Address Information

  Time Server Address : Disabled

   VLAN Name    MAC Address          IP Address
  ------------  -------------------  -------------------
  DEFAULT_VLAN  0001e7-a09900        10.28.227.101
  VLAN-22       0001e7-a09901        Disabled
  VLAN-33       0001e7-a09902        Disabled


 Actions->   Back     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure B-3. Example of Management Address Information with VLANs Configured

This screen displays addresses that are important for management of the switch. If multiple VLANs are *not* configured, this screen displays a single IP address for the entire switch. See the online Help for details.

CLI Access

Syntax: show management

Module Information

Use this feature to determine which slots have modules installed and which type(s) of modules are installed.

Menu: Displaying Port Status

From the Main Menu, select:

1. Status and Counters . . .
3. Module Information

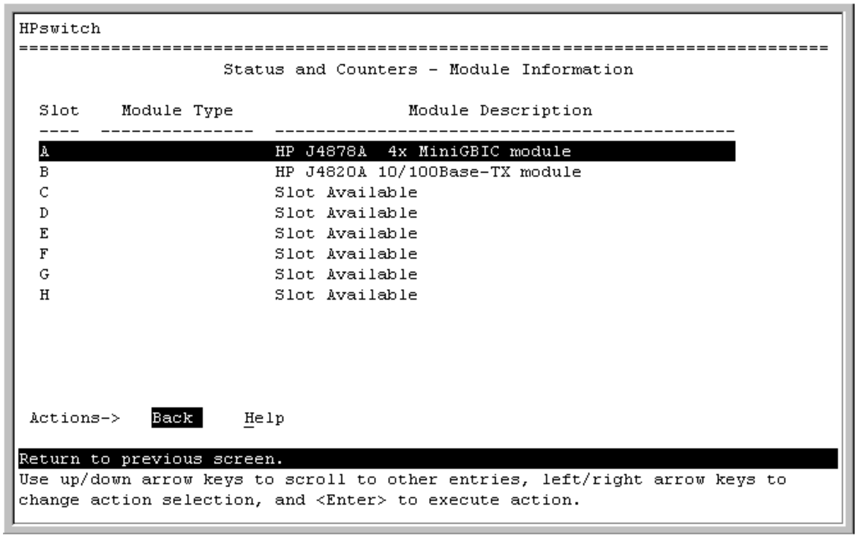

```
HPswitch
================================================================================
                    Status and Counters - Module Information

  Slot    Module Type                 Module Description
  ----  ---------------  ---------------------------------------------
  A                      HP J4878A  4x MiniGBIC module
  B                      HP J4820A 10/100Base-TX module
  C                      Slot Available
  D                      Slot Available
  E                      Slot Available
  F                      Slot Available
  G                      Slot Available
  H                      Slot Available






 Actions->   Back    Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure B-4. Example of Module Information in the Menu Interface

CLI Access

Syntax: show module

Port Status

The web browser interface and the console interface show the same port status data.

Menu: Displaying Port Status

From the Main Menu, select:

1. Status and Counters . . .
4. Port Status

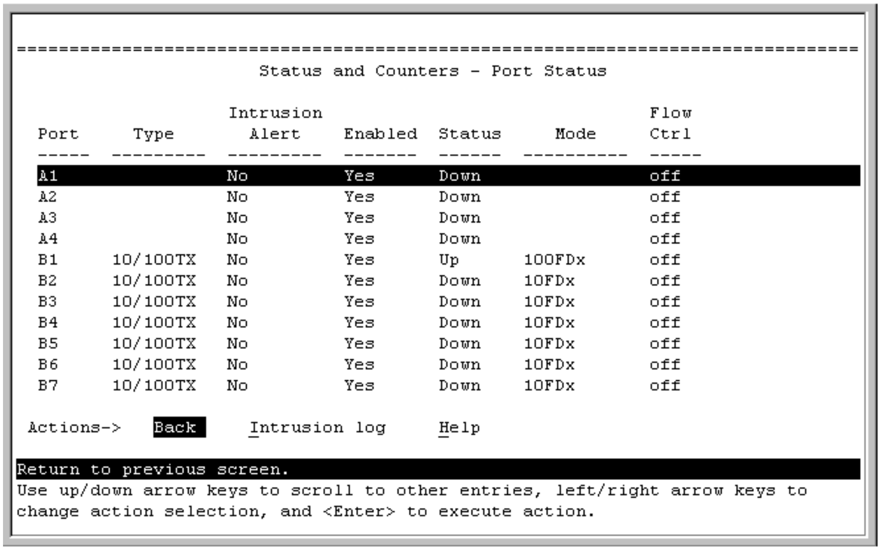

```
=================================================================================
                      Status and Counters - Port Status

                        Intrusion                                  Flow
  Port     Type           Alert    Enabled  Status     Mode        Ctrl
  -----  ---------    ---------    -------  ------  ----------     -----
  A1                    No          Yes      Down                  off
  A2                    No          Yes      Down                  off
  A3                    No          Yes      Down                  off
  A4                    No          Yes      Down                  off
  B1     10/100TX       No          Yes      Up      100FDx        off
  B2     10/100TX       No          Yes      Down    10FDx         off
  B3     10/100TX       No          Yes      Down    10FDx         off
  B4     10/100TX       No          Yes      Down    10FDx         off
  B5     10/100TX       No          Yes      Down    10FDx         off
  B6     10/100TX       No          Yes      Down    10FDx         off
  B7     10/100TX       No          Yes      Down    10FDx         off

 Actions->   Back      Intrusion log     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure B-5. Example of Port Status on the Menu Interface

CLI Access

Syntax: show interfaces brief

Web Access

1. Click on the **Status** tab.
2. Click on Port Status.

Viewing Port and Trunk Group Statistics and Flow Control Status

Feature	Default	Menu	CLI	Web
viewing port and trunk statistics for all ports, and flow control status	n/a	page B-10	page B-11	page B-11
viewing a detailed summary for a particular port or trunk	n/a	page B-10	page B-11	page B-11
resetting counters	n/a	page B-10	page B-11	page B-11

These features enable you to determine the traffic patterns for each port since the last reboot or reset of the switch. You can display:

- A general report of traffic on all LAN ports and trunk groups in the switch, along with the per-port flow control status (On or Off).
- A detailed summary of traffic on a selected port or trunk group.

You can also reset the counters for a specific port.

The menu interface and the web browser interface provide a dynamic display of counters summarizing the traffic on each port. The CLI lets you see a static “snapshot” of port or trunk group statistics at a particular moment.

As mentioned above, rebooting or resetting the switch resets the counters to zero. You can also reset the counters to zero for the current session. This is useful for troubleshooting. See the “Note On Reset”, below.

Note on Reset

The **Reset** action resets the counter display to zero for the current session, but does not affect the cumulative values in the actual hardware counters. (In compliance with the SNMP standard, the values in the hardware counters are not reset to zero unless you reboot the switch.) Thus, using the **Reset** action resets the displayed counters to zero for the current session only. Exiting from the console session and starting a new session restores the counter displays to the accumulated values in the hardware counters.

Menu Access to Port and Trunk Statistics

To access this screen from the Main Menu, select:

1. Status and Counters . . .

4. Port Counters

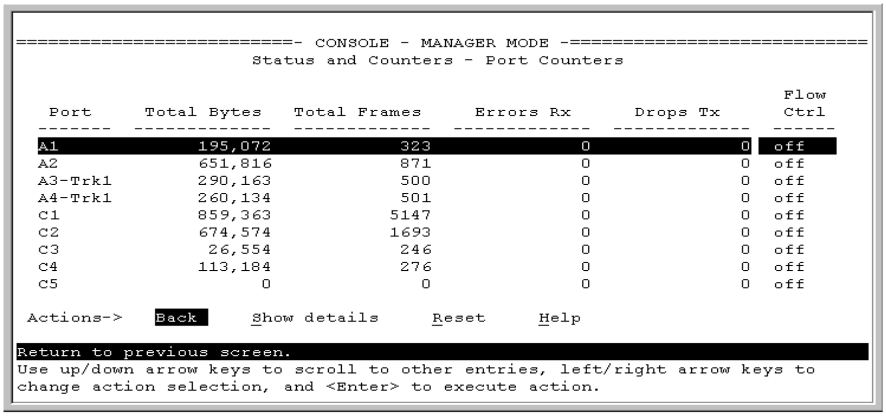

```
==========================- CONSOLE - MANAGER MODE -============================
                     Status and Counters - Port Counters

                                                                        Flow
  Port      Total Bytes    Total Frames     Errors Rx      Drops Tx      Ctrl
 -------   -------------  -------------   -------------  -------------  ------
 A1             195,072            323               0              0   off
 A2             651,816            871               0              0   off
 A3-Trk1        290,163            500               0              0   off
 A4-Trk1        260,134            501               0              0   off
 C1             859,363           5147               0              0   off
 C2             674,574           1693               0              0   off
 C3              26,554            246               0              0   off
 C4             113,184            276               0              0   off
 C5                   0              0               0              0   off

 Actions->   Back      Show details     Reset     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure B-6. Example of Port Counters on the Menu Interface

To view details about the traffic on a particular port, use the [↓] key to highlight that port number, then select **Show Details**. For example, selecting port A2 displays a screen similar to figure B-B-7, below.

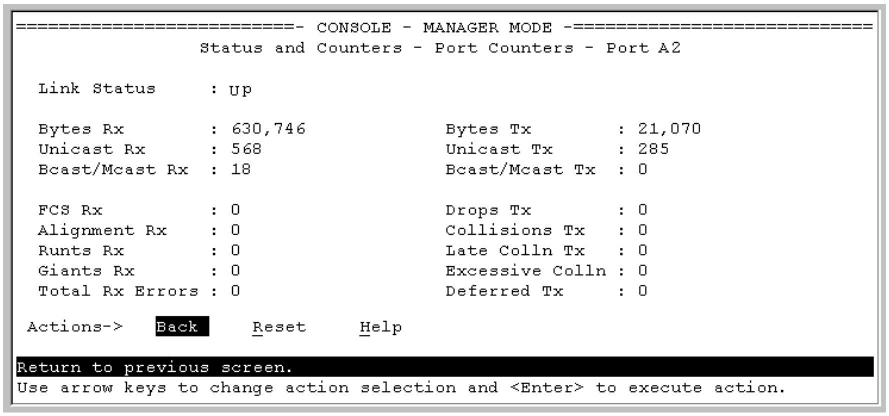

```
==========================- CONSOLE - MANAGER MODE -============================
                 Status and Counters - Port Counters - Port A2

  Link Status     : Up

  Bytes Rx        : 630,746             Bytes Tx        : 21,070
  Unicast Rx      : 568                 Unicast Tx      : 285
  Bcast/Mcast Rx  : 18                  Bcast/Mcast Tx  : 0

  FCS Rx          : 0                   Drops Tx        : 0
  Alignment Rx    : 0                   Collisions Tx   : 0
  Runts Rx        : 0                   Late Colln Tx   : 0
  Giants Rx       : 0                   Excessive Colln : 0
  Total Rx Errors : 0                   Deferred Tx     : 0

 Actions->   Back     Reset     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure B-7. Example of the Display for Show details on a Selected Port

This screen also includes the **Reset** action for the current session. (See the “Note on Reset” on page B-9.)

CLI Access To Port and Trunk Group Statistics

To Display the Port Counter Summary Report. This command provides an overview of port activity for all ports on the switch.

Syntax: show interfaces

To Display a Detailed Traffic Summary for Specific Ports. This command provides traffic details for the port(s) you specify.

Syntax: show interfaces [ethernet] <*port-list*>

To Reset the Port Counters for a Specific Port. This command resets the counters for the specified ports to zero for the current session. (See the “Note on Reset” on page B-9.)

Syntax: clear statistics <[ethernet] *port-list*>

Web Browser Access To View Port and Trunk Group Statistics

1. Click on the **Status** tab.
2. Click on [Port Counters].
3. To reset the counters for a specific port, click anywhere in the row for that port, then click on [Refresh].

Viewing the Switch’s MAC Address Tables

Feature	Default	Menu	CLI	Web
viewing MAC addresses on all ports on a specific VLAN	n/a	page B-12	page B-14	—
viewing MAC addresses on a specific port	n/a	page B-13	page B-14	—
searching for a MAC address	n/a	page B-13	page B-15	—

These features help you to view:

- The MAC addresses that the switch has learned from network devices attached to the switch
- The port on which each MAC address was learned

Menu Access to the MAC Address Views and Searches

Per-VLAN MAC-Address Viewing and Searching. This feature lets you determine which switch port on a selected VLAN is being used to communicate with a specific device on the network. The per-VLAN listing includes:

- The MAC addresses that the switch has learned from network devices attached to the switch
- The port on which each MAC address was learned

1. From the Main Menu, select:

   1. Status and Counters
   5. VLAN Address Table

2. The switch then prompts you to select a VLAN.

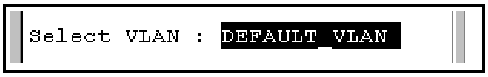

```
Select VLAN : DEFAULT_VLAN
```

3. Use the Space bar to select the VLAN you want, then press [Enter]. The switch then displays the MAC address table for that VLAN:

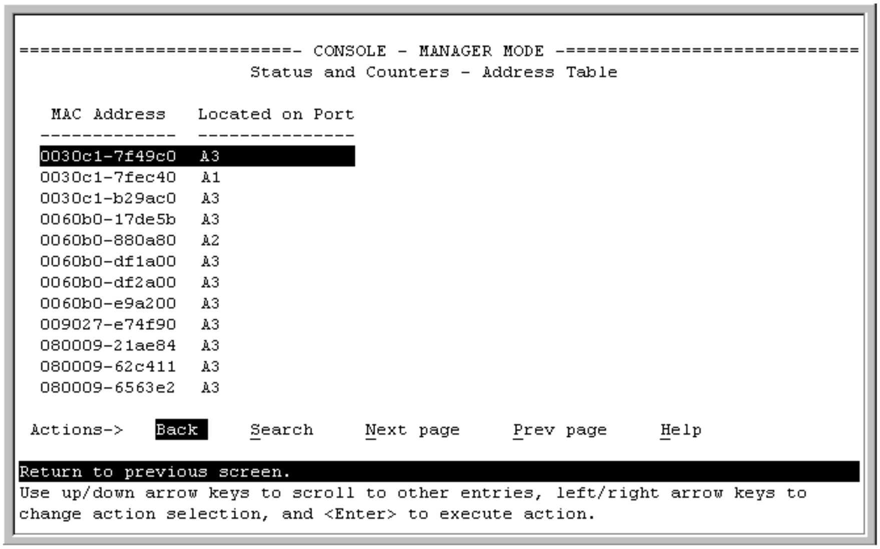

```
==========================- CONSOLE - MANAGER MODE -============================
                        Status and Counters - Address Table

   MAC Address   Located on Port
  -------------  ---------------
  0030c1-7f49c0  A3
  0030c1-7fec40  A1
  0030c1-b29ac0  A3
  0060b0-17de5b  A3
  0060b0-880a80  A2
  0060b0-df1a00  A3
  0060b0-df2a00  A3
  0060b0-e9a200  A3
  009027-e74f90  A3
  080009-21ae84  A3
  080009-62c411  A3
  080009-6563e2  A3

 Actions->   Back     Search     Next page     Prev page     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure B-8. Example of the Address Table

To page through the listing, use **Next page** and **Prev page**.

Finding the Port Connection for a Specific Device on a VLAN. This feature uses a device's MAC address that you enter to identify the port used by that device.

1. Proceeding from figure B-B-8, press [S] (for **Search**), to display the following prompt:

   ```
   Enter MAC address: _
   ```

2. Type the MAC address you want to locate and press [Enter]. The address and port number are highlighted if found. If the switch does not find the MAC address on the currently selected VLAN, it leaves the MAC address listing empty.

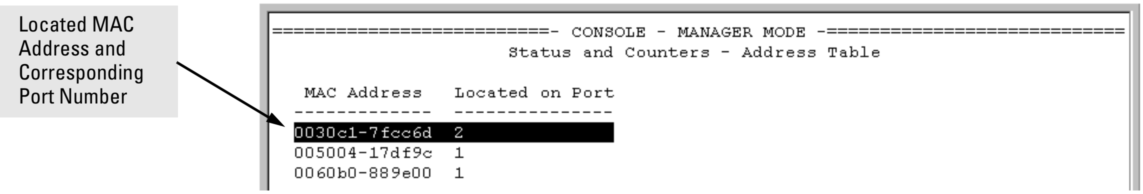

Located MAC Address and Corresponding Port Number

```
==========================- CONSOLE - MANAGER MODE -============================
                        Status and Counters - Address Table

  MAC Address    Located on Port
  -------------  ---------------
 0030c1-7fcc6d  2
 005004-17df9c  1
 0060b0-889e00  1
```

Figure B-9. Example of Menu Indicating Located MAC Address

3. Press [P] (for **Prev page**) to return to the full address table listing.

Port-Level MAC Address Viewing and Searching. This feature displays and searches for MAC addresses on the specified port instead of for all ports on the switch.

1. From the Main Menu, select:

   1. Status and Counters
   7. Port Address Table

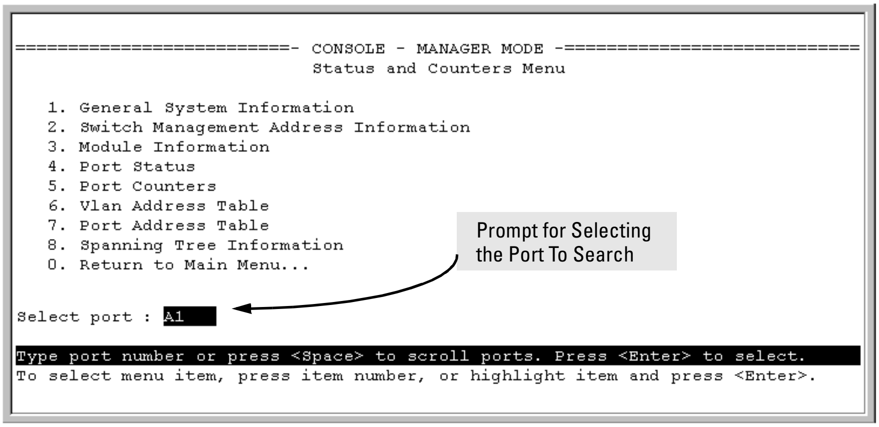

```
=========================- CONSOLE - MANAGER MODE -============================
                            Status and Counters Menu

   1. General System Information
   2. Switch Management Address Information
   3. Module Information
   4. Port Status
   5. Port Counters
   6. Vlan Address Table
   7. Port Address Table
   8. Spanning Tree Information
   0. Return to Main Menu...

Select port : A1

Type port number or press <Space> to scroll ports. Press <Enter> to select.
To select menu item, press item number, or highlight item and press <Enter>.
```

Prompt for Selecting the Port To Search

Figure B-10. Listing MAC Addresses for a Specific Port

2. Use the Space bar to select the port you want to list or search for MAC addresses, then press [Enter] to list the MAC addresses detected on that port.

Determining Whether a Specific Device Is Connected to the Selected Port. Proceeding from step 2, above:

1. Press [S] (for **S**earch), to display the following prompt:

   `Enter MAC address: _`

2. Type the MAC address you want to locate and press [Enter]. The address is highlighted if found. If the switch does not find the address, it leaves the MAC address listing empty.

3. Press [P] (for **Prev page**) to return to the previous per-port listing.

CLI Access for MAC Address Views and Searches

Syntax: show mac-address
  [vlan <*vlan-id*>]
  [ethernet] <*port-list*>]
  [<mac-addr>]

To List All Learned MAC Addresses on the Switch, with The Port Number on Which Each MAC Address Was Learned.

```
HPswitch> show mac-address
```

To List All Learned MAC Addresses on one or more ports, with Their

Corresponding Port Numbers. For example, to list the learned MAC address on ports A1 through A4 and port A6:

```
HPswitch> show mac-address a1-a4,a6
```

To List All Learned MAC Addresses on a VLAN, with Their Port Numbers. This command lists the MAC addresses associated with the ports for a given VLAN. For example:

```
HPswitch> show mac-address vlan 100
```

Note

The Series 5300XL switches operates with a multiple forwarding database architecture. For more on this topic, refer to “Duplicate MAC Addresses on Different Switches” on page C-10

To Find the Port On Which the Switch Learned a Specific MAC Address. For example, to find the port on which the switch learns a MAC address of 080009-21ae84:

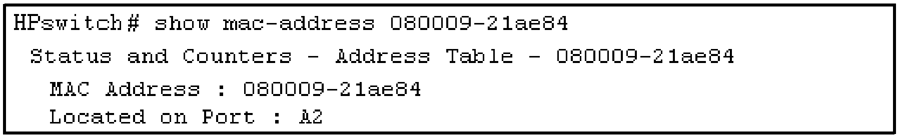

```
HPswitch# show mac-address 080009-21ae84
 Status and Counters - Address Table - 080009-21ae84
  MAC Address : 080009-21ae84
  Located on Port : A2
```

Spanning Tree Protocol (STP) Information

Menu Access to STP Data

From the Main Menu, select:

1. Status and Counters . . .
8. Spanning Tree Information

STP must be enabled on the switch to display the following data:

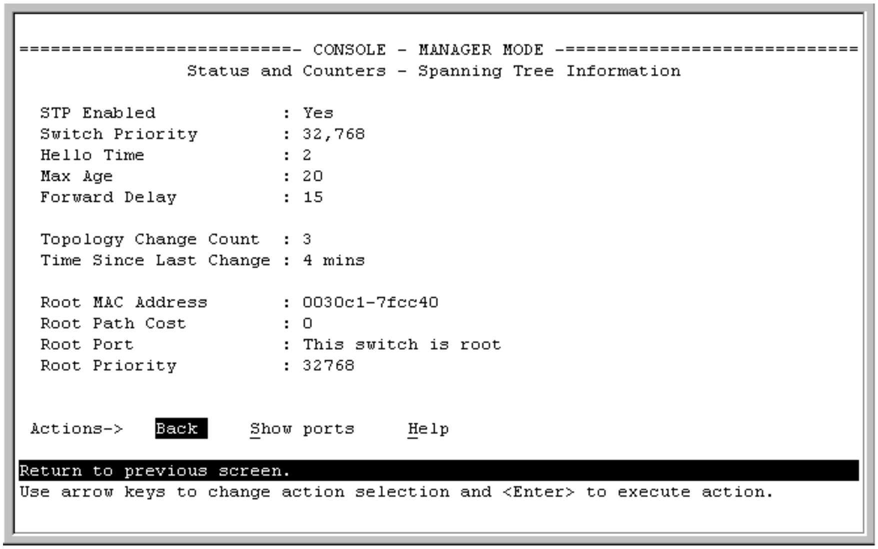

```
==========================- CONSOLE - MANAGER MODE -============================
                Status and Counters - Spanning Tree Information

  STP Enabled            : Yes
  Switch Priority        : 32,768
  Hello Time             : 2
  Max Age                : 20
  Forward Delay          : 15

  Topology Change Count  : 3
  Time Since Last Change : 4 mins

  Root MAC Address       : 0030c1-7fcc40
  Root Path Cost         : 0
  Root Port              : This switch is root
  Root Priority          : 32768


 Actions->   Back     Show ports     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure B-11. Example of Spanning Tree Information

Use this screen to determine current switch-level STP parameter settings and statistics.

You can use the **Show ports** action at the bottom of the screen to display port-level information and parameter settings for each port in the switch (including port type, cost, priority, operating state, and designated bridge) as shown in figure B-B-12.

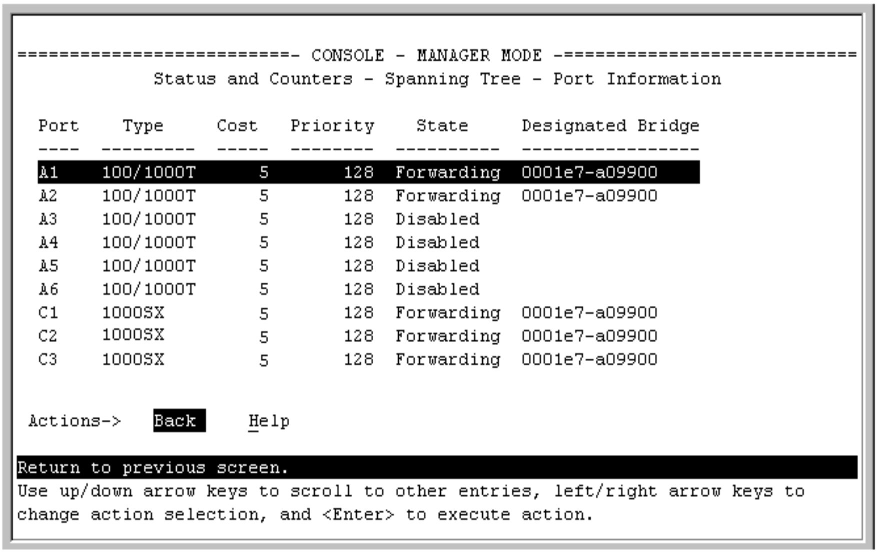

```
==========================- CONSOLE - MANAGER MODE -============================
              Status and Counters - Spanning Tree - Port Information

  Port    Type      Cost   Priority    State     Designated Bridge
  ----  ---------  -----  --------  ----------  -----------------
  A1    100/1000T      5       128  Forwarding  0001e7-a09900
  A2    100/1000T      5       128  Forwarding  0001e7-a09900
  A3    100/1000T      5       128  Disabled
  A4    100/1000T      5       128  Disabled
  A5    100/1000T      5       128  Disabled
  A6    100/1000T      5       128  Disabled
  C1    1000SX         5       128  Forwarding  0001e7-a09900
  C2    1000SX         5       128  Forwarding  0001e7-a09900
  C3    1000SX         5       128  Forwarding  0001e7-a09900


 Actions->    Back     Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure B-12. Example of STP Port Information

CLI Access to STP Data

This option lists the STP configuration, root data, and per-port data (cost, priority, state, and designated bridge).

Syntax: show spanning-tree

```
HPswitch> show spanning-treek
```

Internet Group Management Protocol (IGMP) Status

The switch uses the CLI to display the following IGMP status on a per-VLAN basis:

Show Command	Output
show ip igmp	Global command listing IGMP status for all VLANs configured in the switch: • VLAN ID (VID) and name • Active group addresses per VLAN • Number of report and query packets per group • Querier access port per VLAN
show ip igmp <*vlan-id*>	Per-VLAN command listing above IGMP status for specified VLAN (VID)
show ip igmp group <*ip-addr*>	Lists the ports currently participating in the specified group, with port type, Access type, Age Timer data and Leave Timer data.

For example, suppose that **show ip igmp** listed an IGMP group address of 224.0.1.22. You could get additional data on that group by executing the following:

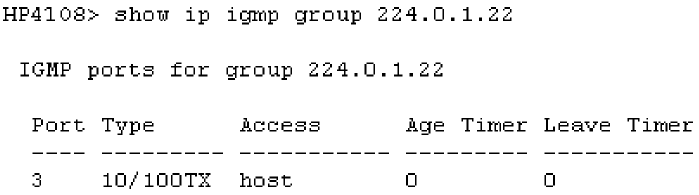

```
HP4108> show ip igmp group 224.0.1.22

 IGMP ports for group 224.0.1.22

  Port Type      Access      Age Timer Leave Timer
  ---- --------- ----------- --------- -----------
  3    10/100TX  host        0         0
```

Figure B-13. Example of IGMP Group Data

VLAN Information

The switch uses the CLI to display the following VLAN status:

Show Command	Output
show vlan	Lists: • Maximum number of VLANs to support • Existing VLANs • Status (static or dynamic) • Primary VLAN
show vlan <*vlan-id*>	For the specified VLAN, lists: • Name, VID, and status (static/dynamic) • Per-Port mode (tagged, untagged, forbid, no/auto) • "Unknown VLAN" setting (Learn, Block, Disable) • Port status (up/down)

For example, suppose that your switch has the following VLANs:

Ports	VLAN	VID
A1 - A12	DEFAULT_VLAN	1
A1, A2	VLAN-33	33
A3, A4	VLAN-44	44

The next three figures show how you could list data on the above VLANs.

Listing the VLAN ID (VID) and Status for ALL VLANs in the Switch.

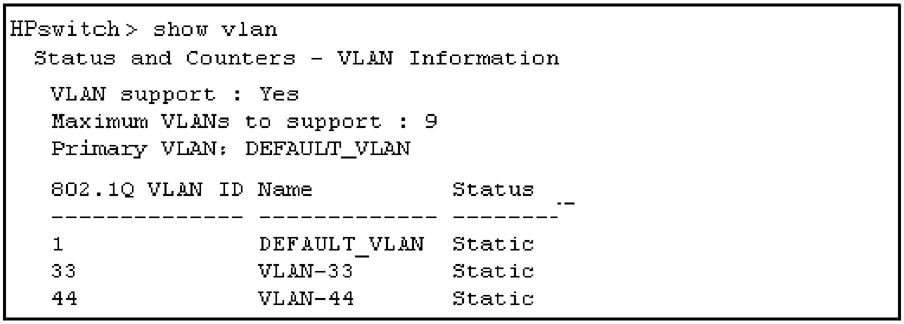

```
HPswitch> show vlan
  Status and Counters - VLAN Information

   VLAN support : Yes
   Maximum VLANs to support : 9
   Primary VLAN: DEFAULT_VLAN

   802.1Q VLAN ID Name          Status
   -------------- ------------- --------
   1              DEFAULT_VLAN  Static
   33             VLAN-33       Static
   44             VLAN-44       Static
```

Figure B-14. Example of VLAN Listing for the Entire Switch

Listing the VLAN ID (VID) and Status for Specific Ports.

Because ports A1 and A2 are not members of VLAN-44, it does not appear in this listing.

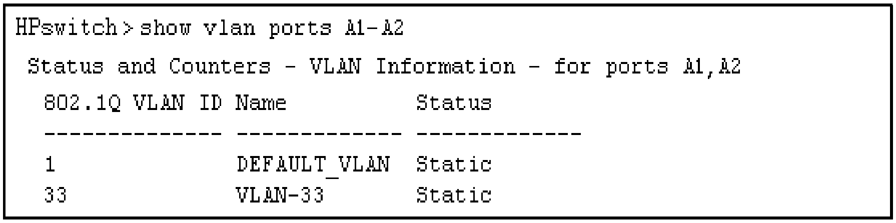

```
HPswitch> show vlan ports A1-A2

 Status and Counters - VLAN Information - for ports A1,A2

  802.1Q VLAN ID Name          Status
  -------------- ------------- -------------
  1              DEFAULT_VLAN  Static
  33             VLAN-33       Static
```

Figure B-15. Example of VLAN Listing for Specific Ports

Listing Individual VLAN Status.

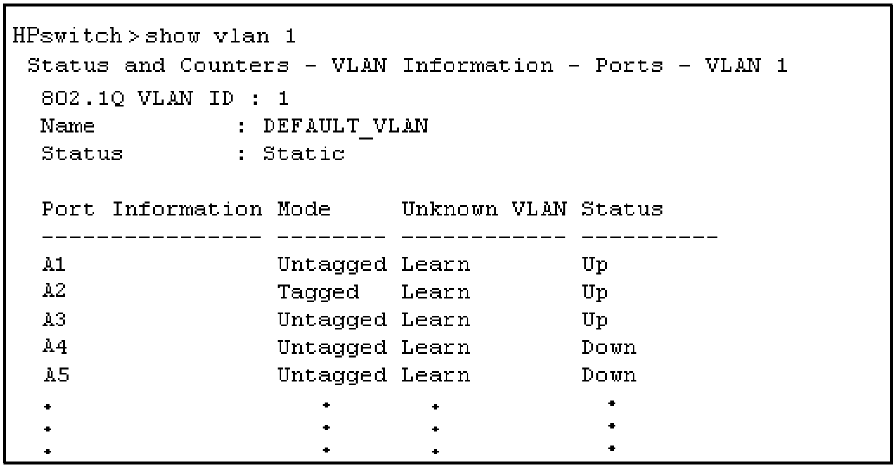

```
HPswitch> show vlan 1
 Status and Counters - VLAN Information - Ports - VLAN 1
  802.1Q VLAN ID : 1
  Name           : DEFAULT_VLAN
  Status         : Static

  Port Information Mode      Unknown VLAN Status
  ---------------- --------- ------------ ----------
  A1               Untagged  Learn        Up
  A2               Tagged    Learn        Up
  A3               Untagged  Learn        Up
  A4               Untagged  Learn        Down
  A5               Untagged  Learn        Down
  .                    .        .            .
  .                    .        .            .
  .                    .        .            .
```

Figure B-16. Example of Port Listing for an Individual VLAN

Web Browser Interface Status Information

The “home” screen for the web browser interface is the Status Overview screen, as shown below. As the title implies, it provides an overview of the status of the switch, including summary graphs indicating the network utilization on each of the switch ports, symbolic port status indicators, and the Alert Log, which informs you of any problems that may have occurred on the switch.

For more information on this screen, see chapter 4, “Using the HP Web Browser Interface”.

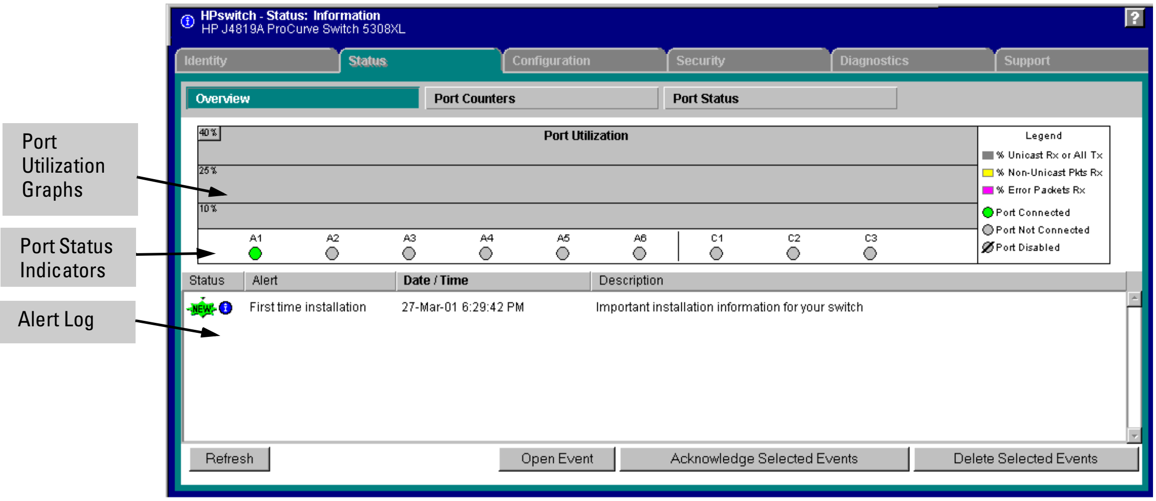

Status	Alert	Date / Time	Description
NEW	First time installation	27-Mar-01 6:29:42 PM	Important installation information for your switch

Figure B-17. Example of a Web Browser Interface Status Overview Screen

Interface Monitoring Features

Port Monitoring Features

Feature	Default	Menu	CLI	Web
display monitoring configuration	disabled	page B-23	page B-25	page B-28
configure the monitor port(s)	ports: none	page B-23	page B-26	page B-28
selecting or removing ports	none selected	page B-23	page B-26	page B-28

You can designate monitoring of inbound traffic on:

- Ports and static trunks: Allows monitoring of individual ports, groups of contiguous ports, and port trunks.
- Static VLANs: Allows traffic monitoring on one static VLAN.
- Meshed ports: Allows traffic monitoring on all ports configured for meshing on the switch.

The switch monitors network activity by copying all traffic inbound on the specified interfaces to the designated monitoring port, to which a network analyzer can be attached.

Note

VLANs, a switch mesh, and port trunks cannot be used as a monitoring port.

It is possible, when monitoring multiple interfaces in networks with high traffic levels, to copy more traffic to a monitor port than the link can support. In this case, some packets may not be copied to the monitor port.

Menu: Configuring Port and Static Trunk Monitoring

This procedure describes configuring the switch for monitoring when monitoring is disabled. (If monitoring has already been enabled, the screens will appear differently than shown in this procedure.)

1. From the Console Main Menu, Select:

   2. Switch Configuration...

   3. Network Monitoring Port

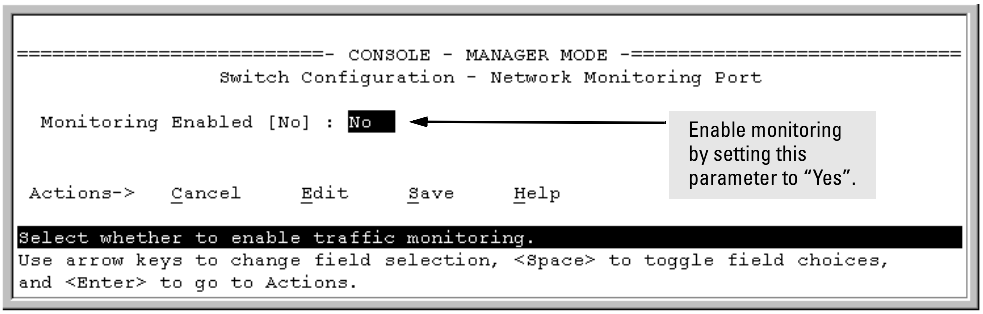

```
==========================- CONSOLE - MANAGER MODE -============================
                  Switch Configuration - Network Monitoring Port

  Monitoring Enabled [No] : No


 Actions->   Cancel     Edit     Save     Help

Select whether to enable traffic monitoring.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Enable monitoring by setting this parameter to “Yes”.

Figure B-18. The Default Network Monitoring Configuration Screen

2. In the Actions menu, press [E] (for `Edit`).
3. If monitoring is currently disabled (the default) then enable it by pressing the Space bar (or [Y]) to select `Yes`.
4. Press the down arrow key to display a screen similar to the following and move the cursor to the **Monitoring Port** parameter.

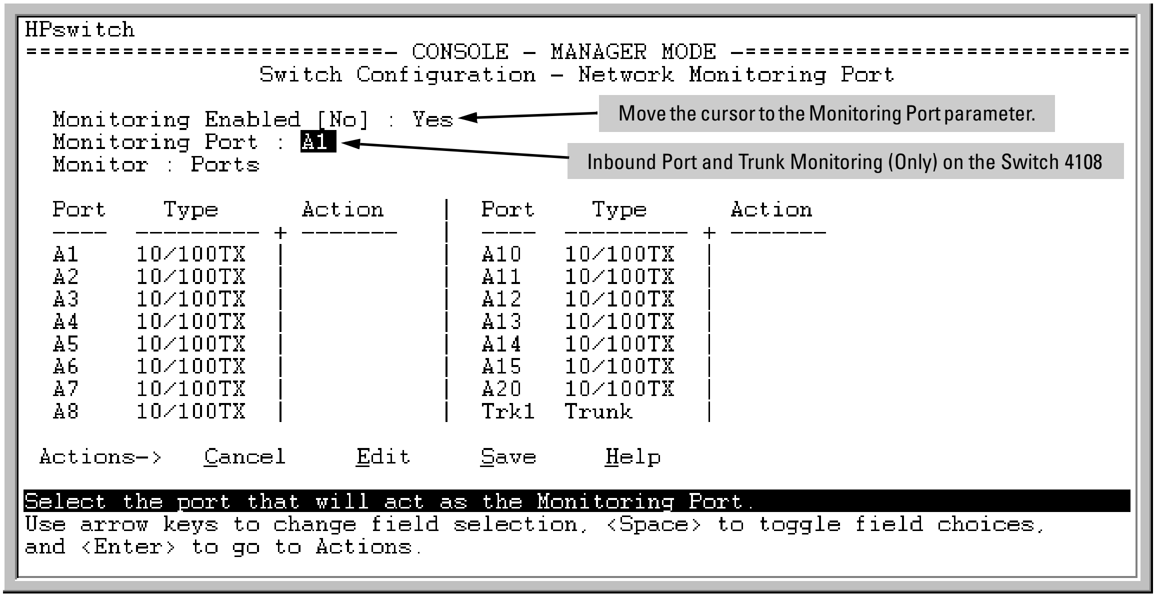

```
HPswitch
==========================- CONSOLE - MANAGER MODE -============================
                  Switch Configuration - Network Monitoring Port

  Monitoring Enabled [No] : Yes
  Monitoring Port : A1
  Monitor : Ports

  Port    Type       Action     |  Port    Type       Action
  ----  ---------- + -------    |  ----  ---------- + -------
  A1    10/100TX   |            |  A10   10/100TX   |
  A2    10/100TX   |            |  A11   10/100TX   |
  A3    10/100TX   |            |  A12   10/100TX   |
  A4    10/100TX   |            |  A13   10/100TX   |
  A5    10/100TX   |            |  A14   10/100TX   |
  A6    10/100TX   |            |  A15   10/100TX   |
  A7    10/100TX   |            |  A20   10/100TX   |
  A8    10/100TX   |            |  Trk1  Trunk      |

 Actions->   Cancel     Edit     Save     Help

Select the port that will act as the Monitoring Port.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Move the cursor to the Monitoring Port parameter.

Inbound Port and Trunk Monitoring (Only) on the Switch 4108

Figure B-19. How To Select a Monitoring Port

5. Use the Space bar to select the port to use for monitoring.
6. Highlight the Monitor field and use the Space bar to select the interfaces to monitor:

   Ports: Use for monitoring ports, static trunks, or the mesh.

   VLAN: Use for monitoring a VLAN.

7. Do one of the following:
   - If you are monitoring ports, static trunks, or the mesh, go to step 8.
   - If you are monitoring a VLAN:
     i. Press [Tab] or the down arrow key to move to the **VLAN** field.

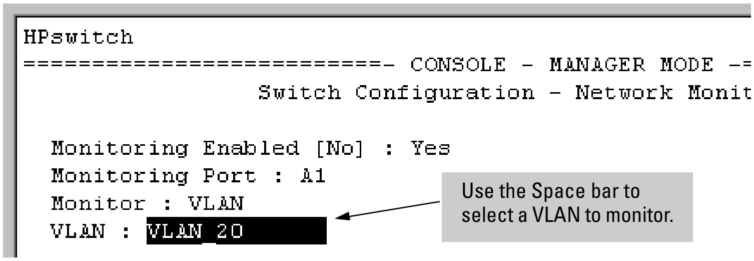

```
HPswitch
==========================- CONSOLE - MANAGER MODE -=
                  Switch Configuration - Network Monit

  Monitoring Enabled [No] : Yes
  Monitoring Port : A1
  Monitor : VLAN
  VLAN : VLAN_20
```

Use the Space bar to
select a VLAN to monitor.

ii. Use the Space bar to select the VLAN you want to monitor.
iii. Go to step 10.

8. Use the down arrow key to move the cursor to the **Action** column for the individual ports and position the cursor at a port you want to monitor.
9. Press the Space bar to select **Monitor** for each port and trunk that you want monitored. (Use the down arrow key to move from one interface to the next in the **Action** column.)
10. When you finish selecting ports to monitor, press [Enter], then press [S] (for **Save**) to save your changes and exit from the screen.
11. Return to the Main Menu.

CLI: Configuring Port and Static Trunk Monitoring

Port and Static Trunk Monitoring Commands Used in This Section

show monitor	below
mirror-port	page B-26
monitor	page B-26

You must use the following configuration sequence to configure port and static trunk monitoring in the CLI:

1. Assign a monitoring (mirror) port.
2. Designate the port(s) and static trunk(s) to monitor.

Displaying the Monitoring Configuration. This command lists the port assigned to receive monitored traffic and the ports and/or trunks being monitored.

Syntax: show monitor

For example, if you assign port A6 as the monitoring port and configure the switch to monitor ports A1 - A3, **show monitor** displays the following:

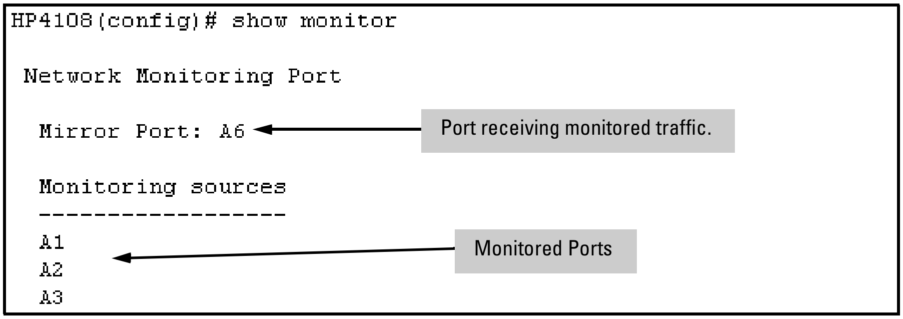

```
HP4108(config)# show monitor

 Network Monitoring Port

  Mirror Port: A6

  Monitoring sources
  ------------------
  A1
  A2
  A3
```

Port receiving monitored traffic.

Monitored Ports

Figure B-20. Example of Monitored Port Listing

Configuring the Monitor Port. This command assigns or removes a monitoring port, and must be executed from the global configuration level. Removing the monitor port disables port monitoring and resets the monitoring parameters to their factory-default settings.

Syntax: [no] mirror-port [< *port-num* >]

For example, to assign port A6 as the monitoring port:

```
HPswitch(config)# mirror-port a6
```

To turn off monitoring:

```
HPswitch(config)# no mirror-port
```

Selecting or Removing Monitoring Source Interfaces. After you configure a monitor port you can use either the global configuration level or the interface context level to select ports, static trunks, meshed ports, or VLANs as monitoring sources. You can also use either level to remove monitoring sources.

Note

Individual ports, static trunks, and meshing can all be monitored at the same time. However, if you configure the switch to monitor a VLAN, all other interfaces are removed from monitoring. Also, you can configure only one VLAN at a time for monitoring.

Syntax: [no] interface [ethernet] < *monitor-list* | vlan < *vlan-id*>> monitor

where: < monitor-list > includes port numbers, static trunk names, and meshing, such as **a4**, **c7**, **b5-b8**, **trk1**, *and* **mesh**.

Elements in the monitor list can include port numbers, static trunk names, and the mesh at the same time.

For example, with a port such as port A6 configured as the monitoring (mirror) port, you would use either of the following commands to select these interfaces for monitoring:

- A1 through A3, and A5
- Trunks 1 and 2
- Meshing

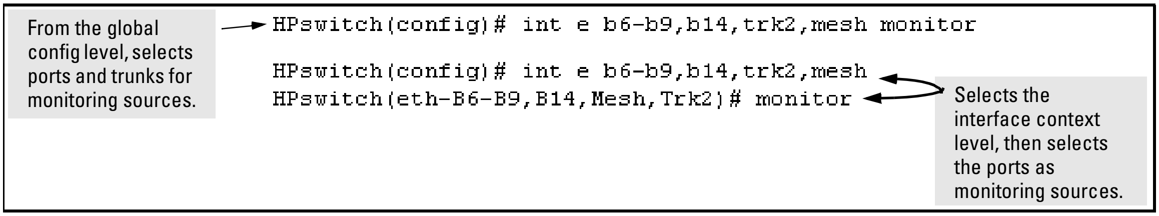

From the global config level, selects ports and trunks for monitoring sources.

```
HPswitch(config)# int e b6-b9,b14,trk2,mesh monitor

HPswitch(config)# int e b6-b9,b14,trk2,mesh
HPswitch(eth-B6-B9,B14,Mesh,Trk2)# monitor
```

Selects the interface context level, then selects the ports as monitoring sources.

Figure B-21. Examples of Selecting Ports and Static Trunks as Monitoring Sources

To monitor a VLAN:

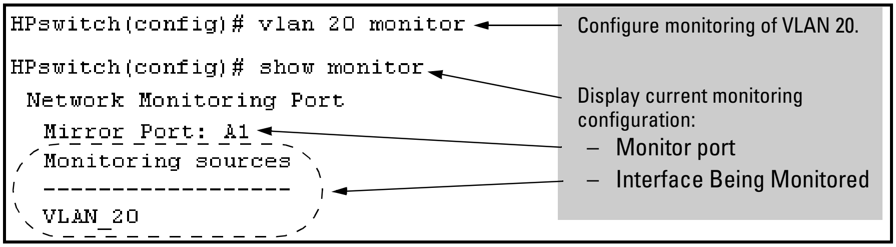

```
HPswitch(config)# vlan 20 monitor

HPswitch(config)# show monitor
 Network Monitoring Port
  Mirror Port: A1
  Monitoring sources
  ------------------
  VLAN_20
```

Configure monitoring of VLAN 20.

Display current monitoring configuration:
– Monitor port
– Interface Being Monitored

Figure B-22. Example of Configuring VLAN Monitoring

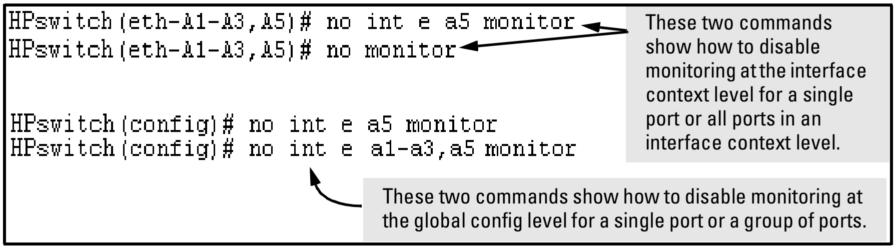

```
HPswitch(eth-A1-A3,A5)# no int e a5 monitor
HPswitch(eth-A1-A3,A5)# no monitor


HPswitch(config)# no int e a5 monitor
HPswitch(config)# no int e  a1-a3,a5 monitor
```

These two commands show how to disable monitoring at the interface context level for a single port or all ports in an interface context level.

These two commands show how to disable monitoring at the global config level for a single port or a group of ports.

Figure B-23. Examples of Removing Ports as Monitoring Sources

Web: Configuring Port Monitoring

To enable port monitoring:

1. Click on the **Configuration** tab.
2. Click on [Monitor Port].
3. To monitor one or more ports.
   a. Click on the radio button for **Monitor Selected Ports**.
   b. Select the port(s) to monitor.
4. Click on [Apply Changes].

To remove port monitoring:

1. Click on the **Monitoring Off** radio button.
2. Click on [Apply Changes].

For web-based Help on how to use the web browser interface screen, click on the [?] button provided on the web browser screen.

Troubleshooting

Contents

Overview

This chapter addresses performance-related network problems that can be caused by topology, switch configuration, and the effects of other devices or their configurations on switch operation. (For switch-specific information on hardware problems indicated by LED behavior, cabling requirements, and other potential hardware-related problems, refer to the installation guide you received with the switch.)

Note

HP periodically places switch software updates on the HP Procurve web site. HP recommends that you check this web site for software updates that may have fixed a problem you are experiencing.

For information on support and warranty provisions, see the Support and Warranty booklet shipped with the switch.

Troubleshooting Approaches

Use these approaches to diagnose switch problems:

- Check the HP Procurve web site for software updates that may have solved your problem: **http://www.hp.com/go/hpprocurve**
- Check the switch LEDs for indications of proper switch operation:
  - Each switch port has a Link LED that should light whenever an active network device is connected to the port.
  - Problems with the switch hardware and software are indicated by flashing the Fault and other switch LEDs.

    See the *Installation Guide* shipped with the switch for a description of the LED behavior and information on using the LEDs for troubleshooting.
- Check the network topology/installation. See the *Installation Guide* shipped with the switch for topology information.
- Check cables for damage, correct type, and proper connections. You should also use a cable tester to check your cables for compliance to the relevant IEEE 802.3 specification. See the *Installation Guide* shipped with the switch for correct cable types and connector pin-outs.
- Use HP TopTools for Hubs & Switches (if installed on your network) to help isolate problems and recommend solutions. HP TopTools is shipped at no extra cost with the switch.
- Use the Port Utilization Graph and Alert Log in the web browser interface included in the switch to help isolate problems. See Chapter 4, "Using the HP Web Browser Interface" for operating information. These tools are available through the web browser interface:
  - Port Utilization Graph
  - Alert Log
  - Port Status and Port Counters screens
  - Diagnostic tools (Link test, Ping test, configuration file browser)
- For help in isolating problems, use the easy-to-access switch console built into the switch or Telnet to the switch console. See chapter 4, "Using the Switch Console Interface" for operating information. These tools are available through the switch console
  - Status and Counters screens
  - Event Log
  - Diagnostics tools (Link test, Ping test, configuration file browser, and advanced user commands)

Browser or Telnet Access Problems

Cannot access the web browser interface:

- Access may be disabled by the **Web Agent Enabled** parameter in the switch console. Check the setting on this parameter by selecting:

  2. Switch Configuration . . .

  1. System Information

- The switch may not have the correct IP address, subnet mask or gateway. Verify by connecting a console to the switch's Console port and selecting:

  2. Switch Configuration . . .

  5. IP Configuration

  Note: If DHCP/Bootp is used to configure the switch, the IP addressing can be verified by selecting:

  1. Status and Counters . . .

  2. Switch Management Address Information

  also check the DHCP/Bootp server configuration to verify correct IP addressing.

- If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" may have expired so that the IP address has changed. For more information on how to "reserve" an IP address, refer to the documentation for the DHCP application that you are using.
- If one or more IP-Authorized managers are configured, the switch allows web browser access only to a device having an authorized IP address. For more information on IP Authorized managers, refer to the *Access Security Guide* for your switch.
- Java™ applets may not be running on the web browser. They are required for the switch web browser interface to operate correctly. See the online Help on your web browser for how to run the Java applets.

Cannot Telnet into the switch console from a station on the network:

- Off subnet management stations can lose Telnet access if you enable routing without first configuring a static (default) route. That is, the switch uses the IP default gateway only while operating as a Layer 2 device. While routing is enabled on the switch, the IP default gateway is not used. You can avoid this problem by using the ip route command to configure a static (default) route before enabling routing. Refer to chapter 16, "IP Routing Features", for more information.
- Telnet access may be disabled by the **Inbound Telnet Enabled** parameter in the System Information screen of the menu interface:

  2. Switch Configuration

  1. System Information
- The switch may not have the correct IP address, subnet mask, or gateway. Verify by connecting a console to the switch's Console port and selecting:

  2. Switch Configuration

  5. IP Configuration

  Note: If DHCP/Bootp is used to configure the switch, see the **Note**, above.
- If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" may have expired so that the IP address has changed. For more information on how to "reserve" an IP address, refer to the documentation for the DHCP application that you are using.
- If one or more IP-Authorized managers are configured, the switch allows inbound telnet access only to a device having an authorized IP address. For more information on IP Authorized managers, refer to the *Access Security Guide* for your switch..

Unusual Network Activity

Network activity that fails to meet accepted norms may indicate a hardware problem with one or more of the network components, possibly including the switch. Such problems can also be caused by a network loop or simply too much traffic for the network as it is currently designed and implemented. Unusual network activity is usually indicated by the LEDs on the front of the switch or measured with the switch console interface or with a network management tool such as the HP TopTools for Hubs & Switches. Refer to the *Installation Guide* you received with the switch for information on using LEDs to identify unusual network activity.

A topology loop can also cause excessive network activity. The event log "FFI" messages can be indicative of this type of problem.

General Problems

The network runs slow; processes fail; users cannot access servers or other devices. Broadcast storms may be occurring in the network. These may be due to redundant links between nodes.

- If you are configuring a port trunk, finish configuring the ports in the trunk before connecting the related cables. Otherwise you may inadvertently create a number of redundant links (i.e. topology loops) that will cause broadcast storms.
- Turn on Spanning Tree Protocol to block redundant links (i.e. topology loops)
- Check for FFI messages in the Event Log.

Duplicate IP Addresses. This is indicated by this Event Log message:

ip: Invalid ARP source: *IP address* **on** *IP address*

where: both instances of *IP address* are the same address, indicating the switch's IP address has been duplicated somewhere on the network.

Duplicate IP Addresses in a DHCP Network. If you use a DHCP server to assign IP addresses in your network and you find a device with a valid IP address that does not appear to communicate properly with the server or other devices, a duplicate IP address may have been issued by the server. This can occur if a client has not released a DHCP-assigned IP address after the intended expiration time and the server "leases" the address to another device.

This can also happen, for example, if the server is first configured to issue IP addresses with an unlimited duration, then is subsequently configured to issue IP addresses that will expire after a limited duration. One solution is to configure "reservations" in the DHCP server for specific IP addresses to be assigned to devices having specific MAC addresses. For more information, refer to the documentation for the DHCP server.

One indication of a duplicate IP address in a DHCP network is this Event Log message:

> **ip: Invalid ARP source:** *IP address* **on** *IP address*
>
> *where:* both instances of *IP address* are the same address, indicating the IP address that has been duplicated somewhere on the network.

The Switch Has Been Configured for DHCP/Bootp Operation, But Has Not Received a DHCP or Bootp Reply. When the switch is first configured for DHCP/Bootp operation, or if it is rebooted with this configuration, it immediately begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch may not immediately receive the desired configuration. After verifying that the server has become accessible to the switch, reboot the switch to re-start the process.

802.1q Prioritization Problems

Ports configured for non-default prioritization (level 1 - 7) are not performing the specified action. If the ports were placed in a trunk group after being configured for non-default prioritization, the priority setting was automatically reset to zero (the default). Ports in a trunk group operate only at the default priority setting.

CDP Problems

The switch does not appear in the CDP Neighbors table of an adjacent CDP Device. This may be due to any of the following:

- Either the port connecting the switch to the adjacent device is not a member of an untagged VLAN or any Untagged VLAN to which the port belongs does not have an IP address.

- If there is more than one physical path between the switch and the other CDP device and STP is running on the switch, then STP will block the redundant link(s). In this case, the switch port on the remaining open link may not be a member of an untagged VLAN, or any untagged VLANs to which the port belongs may not have an IP address.

- The adjacent device's CDP Neighbors table may be full. Refer to the documentation provided for the adjacent CDP device to determine the table's capacity, and then view the device's Neighbors table to determine whether it is full.

One or more CDP neighbors appear intermittently or not at all in the switch's CDP Neighbors table. This may be caused by more than 60 neighboring devices sending CDP packets to the switch. Exceeding the 60-neighbor limit can occur, for example, where multiple neighbors are connected to the switch through non-CDP devices such as many hubs.

The Same CDP Switch or Router Appears on More Than One Port in the CDP Neighbors Table. Where CDP is running, a switch or router that is the STP root transmits outbound CDP packets over all links, including redundant links that STP may be blocking in non-root devices. In this case, the non-root device shows an entry in its CDP Neighbors table for every port on which it receives a CDP packet from the root device. See "Effect of Spanning Tree (STP) On CDP Packet Transmission" on page 10-24.

IGMP-Related Problems

IP Multicast (IGMP) Traffic That Is Directed By IGMP Does Not Reach IGMP Hosts or a Multicast Router Connected to a Port. IGMP must be enabled on the switch and the affected port must be configured for "Auto" or "Forward" operation.

IP Multicast Traffic Floods Out All Ports; IGMP Does Not Appear To Filter Traffic. The IGMP feature does not operate if the switch or VLAN does not have an IP address configured manually or obtained through DHCP/Bootp. To verify whether an IP address is configured for the switch or VLAN, do either of the following:

- **Try Using the Web Browser Interface:** If you can access the web browser interface, then an IP address is configured.
- **Try To Telnet to the Switch Console:** If you can Telnet to the switch, then an IP address is configured.

- **Using the Switch Console Interface:** From the Main Menu, check the Management Address Information screen by clicking on

  1. Status and Counters

  2. Switch Management Address Information

LACP-Related Problems

Unable to enable LACP on a port with the **interface [e] <** ***port-number*** **> lacp** command. In this case, the switch displays the following message:

```
Operation is not allowed for a trunked port.
```

You cannot enable LACP on a port while it is configured as static **Trunk** or **FEC** trunked port. To enable LACP on static-trunked port, first use the **no trunk [e] <** ***port-number*** **>** command to disable the static trunk assignment, then execute **interface [e] <** ***port-number*** **> lacp**.

Caution

Removing a port from a trunk without first disabling the port can create a traffic loop that can slow down or halt your network. Before removing a port from a trunk, HP recommends that you either disable the port or disconnect it from the LAN.

Mesh-Related Problems

Traffic on a dynamic VLAN does not get through the switch mesh .

GVRP enables dynamic VLANs. Ensure that all switches in the mesh have GVRP enabled. (Note that HP Procurve 1600M/2400M/2424M/4000M/8000M switches do not offer GVRP. Thus, if there are any of these switches in the mesh, GVRP must be disabled for any Series 5300XL switches in the mesh.)

The Switch Mesh Does Not Allow An HP Procurve Switch 1600M/

2400M/2424M/4000M/8000M Port To Join the Mesh . One of the Series 5300XL switches in the mesh domain has detected a duplicate MAC address on multiple switches. For example:

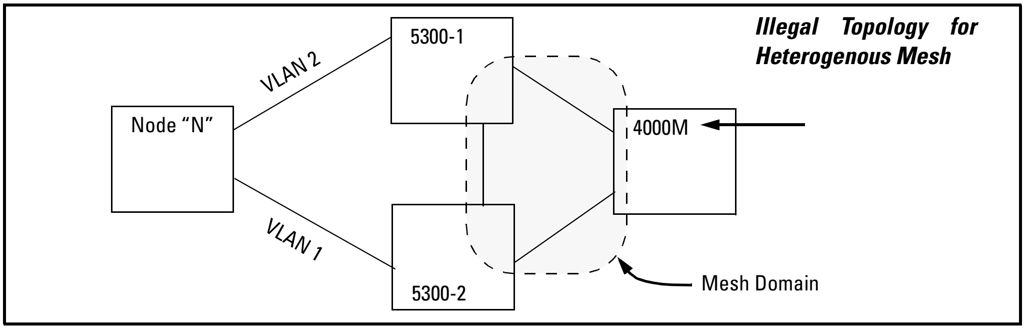

Changing the topology can solve this problem. Also, the duplicate MAC address must age out before the Switch 1600M/2400M/2424M/4000M/8000M port can join the mesh. Refer to "Using a Heterogeneous Switch Mesh" on page 14-7, and to the bulleted item "Compatibility with Older Switches" on page 14-25.

Duplicate MAC Addresses on Different Switches. In a switch mesh that includes any 1600M/2400M/2424M/4000M/8000M switches, duplicate MAC addresses on different switches are not allowed. (The 1600M/2400M/2424M/4000M/8000M switches do not recognize multiple instances of a particular MAC address on different VLANs.) Refer to "The Switch Mesh Does Not Allow An HP Procurve Switch 1600M/2400M/2424M/4000M/8000M Port To Join the Mesh" on page C-9.

Port-Based Access Control (802.1x)-Related Problems

Note To list the 802.1x port-access Event Log messages stored on the switch, use **show log 802**.

See also "Radius-Related Problems" on page C-13.

The switch does not receive a response to RADIUS authentication requests. In this case, the switch will attempt authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use **ping** to ensure that the switch has access to the configured RADIUS servers.
- Verify that the switch is using the correct encryption key (RADIUS secret key) for each server.
- Verify that the switch has the correct IP address for each RADIUS server.
- Ensure that the **radius-server timeout** period is long enough for network conditions.

The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request. If the RADIUS server configuration for authenticating the client includes a VLAN assignment, ensure that the VLAN exists as a static VLAN on the switch. Refer to "How 802.1x Authentication Affects VLAN Operation" in the *Access Security Guide* for your switch.

During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost. If the affected VLAN is configured as untagged on the port, it may be temporarily blocked on that port during an 802.1x session. This is because the switch has temporarily assigned another VLAN as untagged on the port to support the client access, as specified in the response from the RADIUS server. Refer to "How 802.1x Authentication Affects VLAN Operation" in the *Access Security Guide* for your switch.

The switch appears to be properly configured as a supplicant, but cannot gain access to the intended authenticator port on the switch to which it is connected. If **aaa authentication port-access** is configured for Local, ensure that you have entered the local *login* (operator-level) username and password of the authenticator switch into the **identity** and **secret** parameters of the supplicant configuration. If instead, you enter the enable (manager-level) username and password, access will be denied.

The supplicant statistics listing shows multiple ports with the same authenticator MAC address. The link to the authenticator may have been moved from one port to another without the supplicant statistics having been cleared from the first port. Refer to "Note on Supplicant Statistics" in the chapter on Port-Based Access Control in the *Access Security Guide* for your switch.

The show port-access authenticator < ***port-list*** **> command shows one or more ports remain open after they have been configured with control unauthorized.** 802.1x is not active on the switch. After you execute **aaa port-access authenticator active**, all ports configured with **control unauthorized** should be listed as **Closed**.

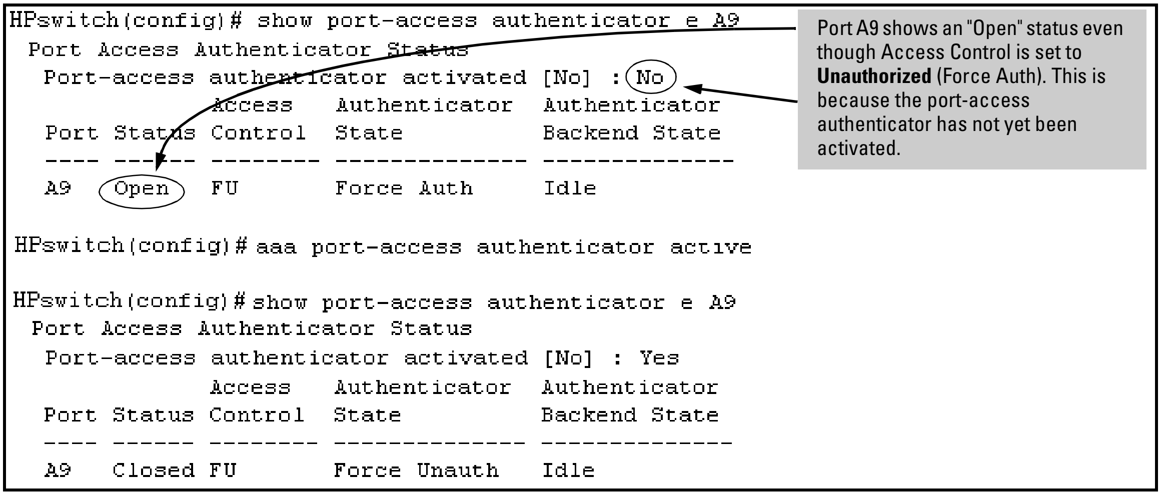

```
HPswitch(config)# show port-access authenticator e A9
 Port Access Authenticator Status
  Port-access authenticator activated [No] : No
               Access   Authenticator  Authenticator
  Port Status Control  State          Backend State
  ---- ------ -------- -------------- --------------
  A9   Open   FU       Force Auth     Idle

HPswitch(config)# aaa port-access authenticator active

HPswitch(config)# show port-access authenticator e A9
 Port Access Authenticator Status
  Port-access authenticator activated [No] : Yes
               Access   Authenticator  Authenticator
  Port Status Control  State          Backend State
  ---- ------ -------- -------------- --------------
  A9   Closed FU       Force Unauth   Idle
```

Port A9 shows an "Open" status even though Access Control is set to **Unauthorized** (Force Auth). This is because the port-access authenticator has not yet been activated.

Figure C-1. Authenticator Ports Remain "Open" Until Activated

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch. Use **show radius** to verify that the encryption key (RADIUS secret key) the switch is using is correct for the server being contacted. If the switch has only a global key configured, then it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, then it overrides the global key and must match the server key.

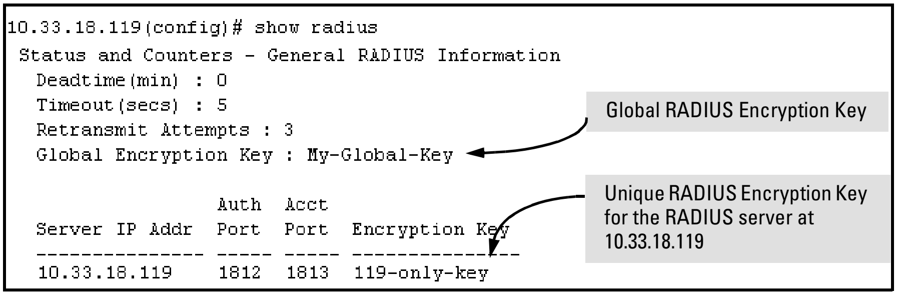

```
10.33.18.119(config)# show radius
 Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key : My-Global-Key

                    Auth  Acct
  Server IP Addr   Port  Port  Encryption Key
  --------------- ----- ----- ---------------
  10.33.18.119     1812  1813  119-only-key
```

Global RADIUS Encryption Key

Unique RADIUS Encryption Key for the RADIUS server at 10.33.18.119

Also, ensure that the switch port used to access the RADIUS server is not blocked by an 802.1x configuration on that port. For example, **show port-access authenticator <** ***port-list*** **>** gives you the status for the specified ports. Also, ensure that other factors, such as port security or any 802.1x configuration on the RADIUS server are not blocking the link.

The authorized MAC address on a port that is configured for both 802.1x and port security either changes or is re-acquired after execution of aaa port-access authenticator < *port-list* > initialize. If the port is force-authorized with **aaa port-access authenticator <port-list> control authorized** command and port security is enabled on the port, then executing **initialize** causes the port to clear the learned address and learn a new address from the first packet it receives after you execute **initialize**.

A trunked port configured for 802.1x is blocked. If you are using RADIUS authentication and the RADIUS server specifies a VLAN for the port, the switch allows authentication, but blocks the port. To eliminate this problem, either remove the port from the trunk or reconfigure the RADIUS server to avoid specifying a VLAN.

Radius-Related Problems

The switch does not receive a response to RADIUS authentication requests. In this case, the switch will attempt authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use **ping** to ensure that the switch has access to the configured RADIUS server.
- Verify that the switch is using the correct encryption key for the designated server.
- Verify that the switch has the correct IP address for the RADIUS server.
- Ensure that the **radius-server timeout** period is long enough for network conditions.
- Verify that the switch is using the same UDP port number as the server.

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch. Use **show radius** to verify that the encryption key the switch is using is correct for the server being contacted. If the switch has only a global key configured, then it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, then it overrides the global key and must match the server key.

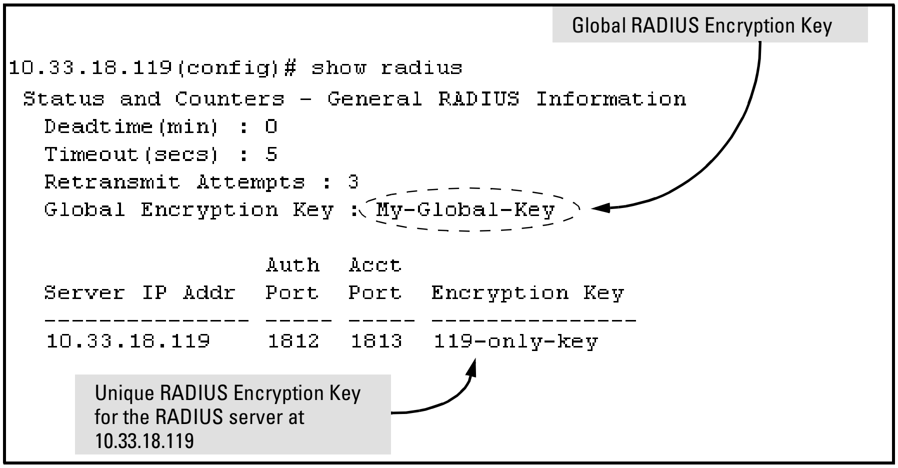

```
10.33.18.119(config)# show radius
 Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key : My-Global-Key

                   Auth  Acct
  Server IP Addr   Port  Port  Encryption Key
  ---------------- ----- ----- ----------------
  10.33.18.119     1812  1813  119-only-key
```

Global RADIUS Encryption Key

Unique RADIUS Encryption Key for the RADIUS server at 10.33.18.119

Figure C-2. Examples of Global and Unique Encryption Keys

Spanning-Tree Protocol (STP) and Fast-Uplink Problems

Caution

If you enable STP, it is recommended that you leave the remainder of the STP parameter settings at their default values until you have had an opportunity to evaluate STP performance in your network. Because incorrect STP settings can adversely affect network performance, you should avoid making changes without having a strong understanding of how STP operates. To learn the details of STP operation, refer to the IEEE 802.1d standard.

Broadcast Storms Appearing in the Network. This can occur when there are physical loops (redundant links) in the topology.Where this exists, you should enable STP on all bridging devices in the topology in order for the loop to be detected.

STP Blocks a Link in a VLAN Even Though There Are No Redundant Links in that VLAN. In 802.1Q-compliant switches such as the Series 5300XL switch, STP blocks redundant physical links even if they are in separate VLANs. A solution is to use only one, multiple-VLAN (tagged) link between the devices. Also, if ports are available, you can improve the bandwidth in this situation by using a port trunk. See “Spanning Tree Operation with VLANs” on page 11-31.

Fast-Uplink Troubleshooting. Some of the problems that can result from incorrect usage of Fast-Uplink STP include temporary loops and generation of duplicate packets.

Problem sources can include:

- Fast-Uplink is configured on a switch that is the STP root device.
- Either the **Hello Time** or the **Max Age** setting (or both) is too long on one or more switches. Return the **Hello Time** and Max Age settings to their default values (2 seconds and 20 seconds, respectively, on a switch).
- A “downlink” port is connected to a switch that is further away (in hop count) from the root device than the switch port on which fast-uplink STP is configured.
- Two edge switches are directly linked to each other with a fast-uplink (Mode = **Uplink**) connection.
- Fast uplink is configured on both ends of a link.
- A switch serving as a backup STP root switch has ports configured for fast-uplink STP and has become the root device due to a failure in the original root device.

SSH-Related Problems

Switch access refused to a client. Even though you have placed the client's public key in a text file and copied the file (using the **copy tftp pub-key-file** command) into the switch, the switch refuses to allow the client to have access. If the source SSH client is an SSHv2 application, the public key may be in the PEM format, which the switch (SSHv1) does not interpret. Check the SSH client application for a utility that can convert the PEM-formatted key into an ASCII-formatted key.

Executing ip ssh does not enable SSH on the switch. The switch does not have a host key. Verify by executing show ip host-public-key. If you see the message

```
ssh cannot be enabled until a host key is configured
(use 'crypto' command).
```

then you need to generate an SSH key pair for the switch. To do so, execute **crypto key generate**.(Refer to "2. Generating the Switch's Public and Private Key Pair" in the SSH chapter of the *Access Security Guide* for your switch.)

Switch does not detect a client's public key that does appear in the switch's public key file (show ip client-public-key). The client's public key entry in the public key file may be preceded by another entry that does not terminate with a new line (CR). In this case, the switch interprets the next sequential key entry as simply a comment attached to the preceding key entry. Where a public key file has more than one entry, ensure that all entries terminate with a new line (CR). While this is optional for the last entry in the file, not adding a new line to the last entry creates an error potential if you either add another key to the file at a later time or change the order of the keys in the file.

An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages.

```
Download failed: overlength key in key file.

Download failed: too many keys in key file.

Download failed: one or more keys is not a valid RSA
public key.
```

The public key file you are trying to download has one of the following problems:

- A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a <CR><LF>.
- There are more than ten public keys in the key file.
- One or more keys in the file is corrupted or is not a valid rsa public key.

Client ceases to respond ("hangs") during connection phase. The switch does not support data compression in an SSH session. Clients will often have compression turned on by default, but will disable it during the negotiation phase. A client which does not recognize the compression-request FAILURE response may fail when attempting to connect. Ensure that compression is turned off before attempting a connection to prevent this problem.

TACACS-Related Problems

Event Log. When troubleshooting TACACS+ operation, check the switch's Event Log for indications of problem areas.

All Users Are Locked Out of Access to the Switch. If the switch is functioning properly, but no username/password pairs result in console or Telnet access to the switch, the problem may be due to how the TACACS+ server and/or the switch are configured. Use one of the following methods to recover:

- Access the TACACS+ server application and adjust or remove the configuration parameters controlling access to the switch.
- If the above method does not work, try eliminating configuration changes in the switch that have not been saved to flash (boot-up configuration) by causing the switch to reboot from the boot-up configuration (which includes only the configuration changes made prior to the last **write memory** command.) If you did not use **write memory** to save the authentication configuration to flash, then pressing the Reset button or cycling the power reboots the switch with the boot-up configuration.
- Disconnect the switch from network access to any TACACS+ servers and then log in to the switch using either Telnet or direct console port access. Because the switch cannot access a TACACS+ server, it will default to local authentication. You can then use the switch's local Operator or Manager username/password pair to log on.

- As a last resort, use the Clear/Reset button combination to reset the switch to its factory default boot-up configuration. Taking this step means you will have to reconfigure the switch to return it to operation in your network.

No Communication Between the Switch and the TACACS+ Server Application. If the switch can access the server device (that is, it can **ping** the server), then a configuration error may be the problem. Some possibilities include:

- The server IP address configured with the switch's tacacs-server host command may not be correct. (Use the switch's **show tacacs-server** command to list the TACACS+ server IP address.)
- The encryption key configured in the server does not match the encryption key configured in the switch (by using the **tacacs-server key** command). Verify the key in the server and compare it to the key configured in the switch. (Use **show tacacs-server** to list the global key. Use **show config** or **show config running** to list any server-specific keys.)
- The accessible TACACS+ servers are not configured to provide service to the switch.

Access Is Denied Even Though the Username/Password Pair Is Correct. Some reasons for denial include the following parameters controlled by your TACACS+ server application:

- The account has expired.
- The access attempt is through a port that is not allowed for the account.
- The time quota for the account has been exhausted.
- The time credit for the account has expired.
- The access attempt is outside of the time frame allowed for the account.
- The allowed number of concurrent logins for the account has been exceeded

For more help, refer to the documentation provided with your TACACS+ server application.

Unknown Users Allowed to Login to the Switch. Your TACACS+ application may be configured to allow access to unknown users by assigning them the privileges included in a *default user* profile. Refer to the documentation provided with your TACACS+ server application.

System Allows Fewer Login Attempts than Specified in the Switch Configuration. Your TACACS+ server application may be configured to allow fewer login attempts than you have configured in the switch with the **aaa authentication num-attempts** command.

TimeP, SNTP, or Gateway Problems

The Switch Cannot Find the Time Server or the Configured Gateway .

TimeP, SNTP, and Gateway access are through the primary VLAN, which in the default configuration is the DEFAULT_VLAN. If the primary VLAN has been moved to another VLAN, it may be disabled or does not have ports assigned to it.

VLAN-Related Problems

Monitor Port. When using the monitor port in a multiple VLAN environment, the switch handles broadcast, multicast, and unicast traffic output from the monitor port as follows:

- If the monitor port is configured for tagged VLAN operation on the same VLAN as the traffic from monitored ports, the traffic output from the monitor port carries the same VLAN tag.
- If the monitor port is configured for untagged VLAN operation on the same VLAN as the traffic from the monitored ports, the traffic output from the monitor port is untagged.
- If the monitor port is not a member of the same VLAN as the traffic from the monitored ports, traffic from the monitored ports does not go out the monitor port.

None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized. If multiple VLANs are being used on ports connecting 802.1Q-compliant devices, inconsistent VLAN IDs may have been assigned to one or more VLANs. For a given VLAN, the same VLAN ID must be used on all connected 802.1Q-compliant devices.

Link Configured for Multiple VLANs Does Not Support Traffic for One

or More VLANs. One or more VLANs may not be properly configured as "Tagged" or "Untagged". A VLAN assigned to a port connecting two 802.1Q-compliant devices must be configured the same on both ports. For example, VLAN_1 and VLAN_2 use the same link between switch "X" and switch "Y".

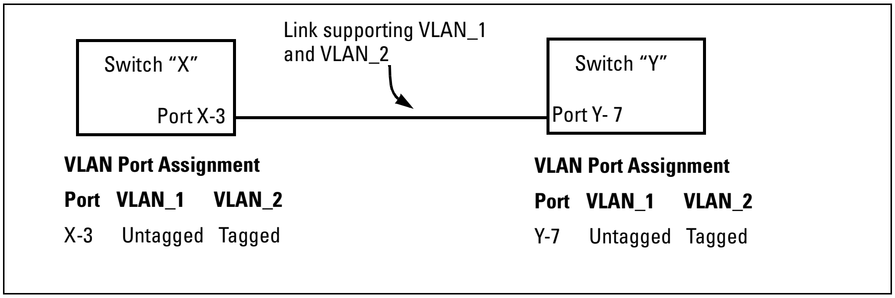

VLAN Port Assignment

Port	VLAN_1	VLAN_2
X-3	Untagged	Tagged

VLAN Port Assignment

Port	VLAN_1	VLAN_2
Y-7	Untagged	Tagged

Figure C-3. Example of Correct VLAN Port Assignments on a Link

1. If VLAN_1 (VID=1) is configured as "Untagged" on port 3 on switch "X", then it must also be configured as "Untagged" on port 7 on switch "Y". Make sure that the VLAN ID (VID) is the same on both switches.
2. Similarly, if VLAN_2 (VID=2) is configured as "Tagged on the link port on switch "A", then it must also be configured as "Tagged" on the link port on switch "B". Make sure that the VLAN ID (VID) is the same on both switches.

Duplicate MAC Addresses Across VLANs. The Series 5300XL switch operates with multiple forwarding databases. Thus, duplicate MAC addresses occurring on different VLANs can appear where a device having one MAC address is a member of more than one 802.1Q VLAN, and the switch port to which the device is linked is using VLANs (instead of STP or trunking) to establish redundant links to another switch. If the other device sends traffic over multiple VLANs, its MAC address will consistently appear in multiple VLANs on the switch port to which it is linked.

Note that attempting to create redundant paths through the use of VLANs will cause problems with some switches. One symptom is that a duplicate MAC address appears in the Port Address Table of one port, and then later appears on another port. While the Series 5300XL switch has multiple forwarding databases, and thus does not have this problem, some switches with a single forwarding database for all VLANs may produce the impression that a connected device is moving among ports because packets with the same MAC address but different VLANs are received on different ports. You can avoid this problem by creating redundant paths using port trunks or spanning tree.

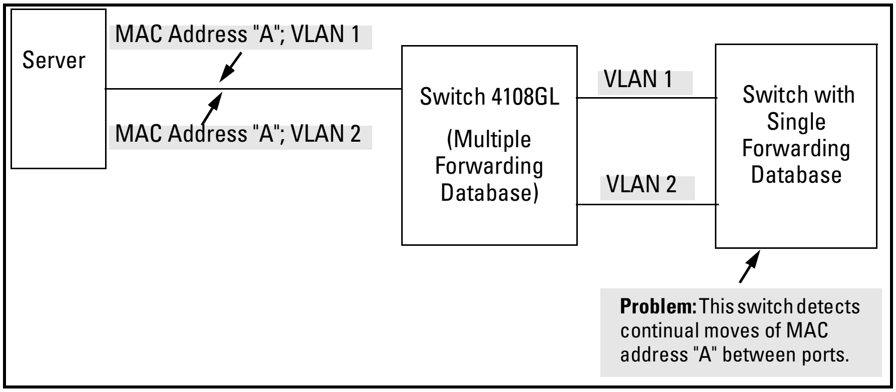

Problem: This switch detects continual moves of MAC address "A" between ports.

Figure C-4. Example of Duplicate MAC Address

Using the Event Log To Identify Problem Sources

The Event Log records operating events as single-line entries listed in chronological order, and serves as a tool for isolating problems. Each Event Log entry is composed of five fields:

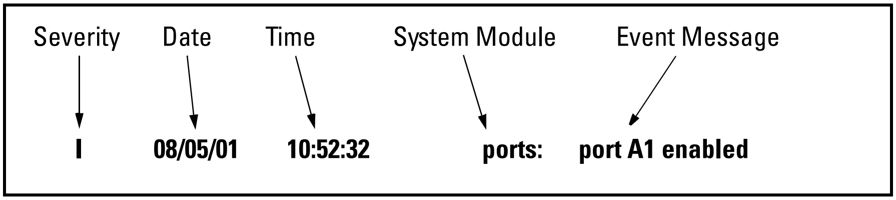

```
I   08/05/01   10:52:32      ports:   port A1 enabled
```

Severity is one of the following codes:

- **I** (information) indicates routine events.
- **W** (warning) indicates that a service has behaved unexpectedly.
- **C** (critical) indicates that a severe switch error has occurred.
- **D** (debug) reserved for HP internal diagnostic information.

Date is the date in *mm/dd/yy* format that the entry was placed in the log.

Time is the time in *hh:mm:ss* format that the entry was placed in the log.

System Module is the internal module (such as “ports” for port manager) that generated the log entry. If VLANs are configured, then a VLAN name also appears for an event that is specific to an individual VLAN. Table 1 on page C-23 lists the individual modules.

Event Message is a brief description of the operating event.

The event log holds up to 1000 lines in chronological order, from the oldest to the newest. Each line consists of one complete event message. Once the log has received 1000 entries, it discards the current oldest line each time a new line is received. The event log window contains 14 log entry lines and can be positioned to any location in the log.

The event log will be *erased* if power to the switch is interrupted.

(The event log is *not* erased by using the **Reboot Switch** command in the Main Menu.)

Table C-1. Event Log System Modules

Module	Event Description
addrMgr	Address table
chassis	switch hardware
bootp	bootp addressing
cdp	
console	Console interface
dhcp	DHCP addressing
download	file transfer
FFI	Find, Fix, and Inform -- available in the console event log and web browser interface alert log
garp	GARP/GVRP
igmp	IP Multicast
ip	IP-related
ipx	Novell Netware
lacp	Dynamic LACP trunks
ldbal	Load-Balance Protocol (meshing)
mgr	Console management
ports	Change in port status; static trunks
snmp	SNMP communications
stp	Spanning Tree
sys, system	Switch management
telnet	Telnet activity
tcp	Transmission control
tftp	File transfer for new OS or config.
timep	Time protocol
vlan	VLAN operations
Xmodem	Xmodem file transfer

Menu: Entering and Navigating in the Event Log

From the Main Menu, select **Event Log**.

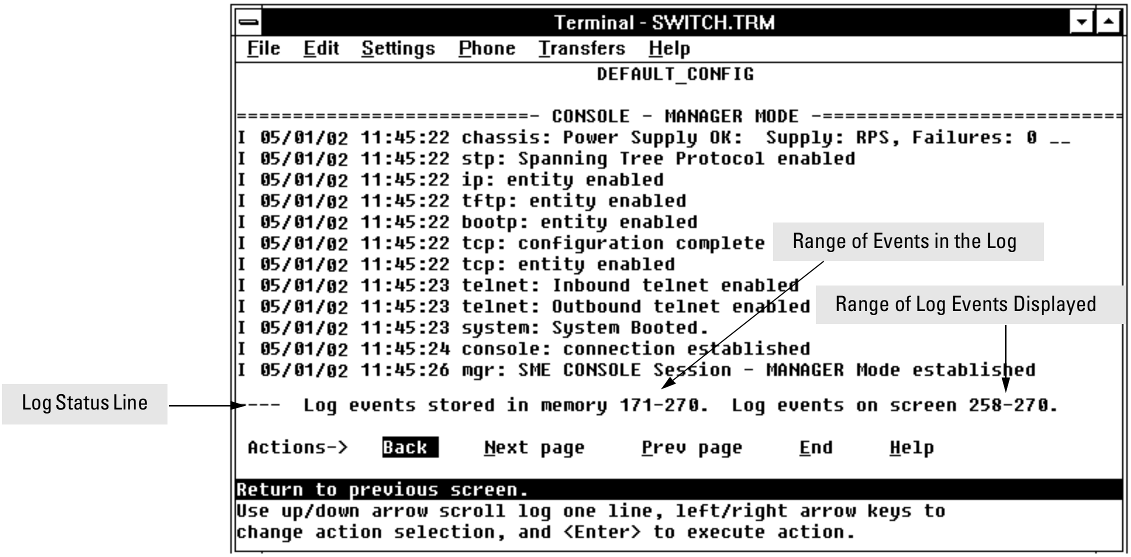

```
                                    DEFAULT_CONFIG

=========================- CONSOLE - MANAGER MODE -===========================
I 05/01/02 11:45:22 chassis: Power Supply OK:  Supply: RPS, Failures: 0 __
I 05/01/02 11:45:22 stp: Spanning Tree Protocol enabled
I 05/01/02 11:45:22 ip: entity enabled
I 05/01/02 11:45:22 tftp: entity enabled
I 05/01/02 11:45:22 bootp: entity enabled
I 05/01/02 11:45:22 tcp: configuration complete
I 05/01/02 11:45:22 tcp: entity enabled
I 05/01/02 11:45:23 telnet: Inbound telnet enabled
I 05/01/02 11:45:23 telnet: Outbound telnet enabled
I 05/01/02 11:45:23 system: System Booted.
I 05/01/02 11:45:24 console: connection established
I 05/01/02 11:45:26 mgr: SME CONSOLE Session - MANAGER Mode established

---  Log events stored in memory 171-270.  Log events on screen 258-270.

 Actions->  Back     Next page     Prev page     End     Help

Return to previous screen.
Use up/down arrow scroll log one line, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Range of Events in the Log

Range of Log Events Displayed

Log Status Line

Figure C-5. Example of an Event Log Display

The *log status line* at the bottom of the display identifies where in the sequence of event messages the display is currently positioned.

To display various portions of the Event Log, either preceding or following the currently visible portion, use either the actions listed at the bottom of the display (**Next page**, **Prev page**, or **End**), or the keys described in the following table:

Table C-2. Event Log Control Keys

Key	Action
[N]	Advance the display by one page (next page).
[P]	Roll back the display by one page (previous page).
[↓]	Advance display by one event (down one line).
[↑]	Roll back display by one event (up one line).
[E]	Advance to the end of the log.
[H]	Display Help for the event log.

CLI:

Using the CLI, you can list

- Events recorded since the last boot of the switch
- All events recorded
- Event entries containing a specific keyword, either since the last boot or all events recorded

Syntax: show logging [-a] [<search-text>]

`HPswitch> show logging`	Lists recorded log messages since last reboot.
`HPswitch> show logging -a`	Lists all recorded log messages, including those before the last reboot.
`HPswitch> show logging -a system`	Lists log messages with “system” in the text or modulename.
`HPswitch> show logging system`	Lists all log messages since the last reboot that have “system” in the text or module name.

Diagnostic Tools

Diagnostic Features

Feature	Default	Menu	CLI	Web
Port Auto negotiation	n/a	n/a	n/a	n/a
Ping Test	n/a	—	page C-29	page C-28
Link Test	n/a	—	page C-29	page C-28
Display Config File	n/a	—	page C-31	page C-31
Admin. and Troubleshooting Commands	n/a	—	page C-33	—
Factory-Default Config	page C-34 (Buttons)	—	page C-34	—
Port Status	n/a	pages B-8 and B-9	pages B-8 and B-9	pages B-8 and B-9

Port Auto-Negotiation

When a link LED does not light (indicating loss of link between two devices), the most common reason is a failure of port auto-negotiation between the connecting ports. If a link LED fails to light when you connect the switch to a port on another device, do the following:

1. Ensure that the switch port and the port on the attached end-node are both set to **Auto** mode.

2. If the attached end-node does not have an **Auto** mode setting, then you must manually configure the switch port to the same setting as the end-node port. See Chapter 9, “Optimizing Traffic Flow with Port Controls, Port Trunking, and Filters”.

Ping and Link Tests

The Ping test and the Link test are point-to-point tests between your switch and another IEEE 802.3-compliant device on your network. These tests can tell you whether the switch is communicating properly with another device.

Note

To respond to a Ping test or a Link test, the device you are trying to reach must be IEEE 802.3-compliant.

Ping Test. This is a test of the path between the switch and another device on the same or another IP network that can respond to IP packets (ICMP Echo Requests).

Link Test. This is a test of the connection between the switch and a designated network device on the same LAN (or VLAN, if configured). During the link test, IEEE 802.2 test packets are sent to the designated network device in the same VLAN or broadcast domain. The remote device must be able to respond with an 802.2 Test Response Packet.

Web: Executing Ping or Link Tests

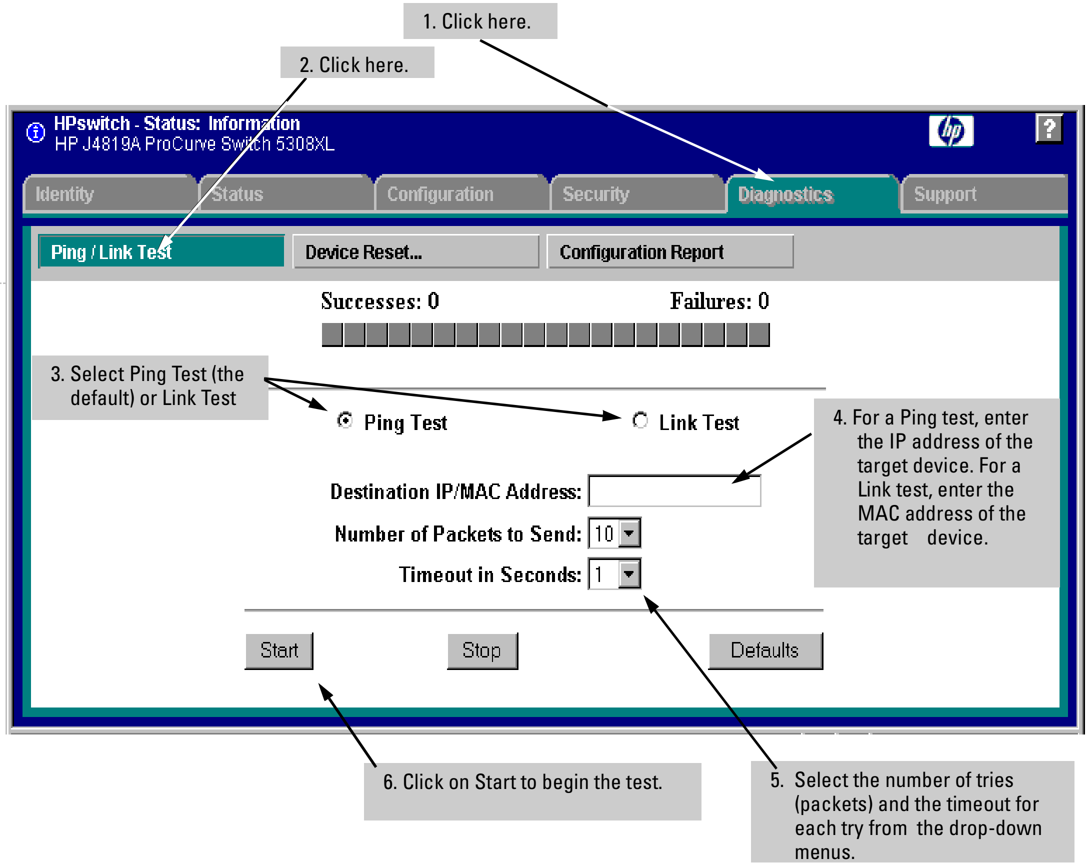

1. Click here.

2. Click here.

HPswitch - Status: Information
HP J4819A ProCurve Switch 5308XL

Successes: 0 Failures: 0

3. Select Ping Test (the default) or Link Test

● Ping Test ○ Link Test

4. For a Ping test, enter the IP address of the target device. For a Link test, enter the MAC address of the target device.

Destination IP/MAC Address:

Number of Packets to Send: 10

Timeout in Seconds: 1

6. Click on Start to begin the test.

5. Select the number of tries (packets) and the timeout for each try from the drop-down menus.

Figure C-6. Link and Ping Test Screen on the Web Browser Interface

Successes indicates the number of Ping or Link packets that successfully completed the most recent test.

Failures indicates the number of Ping or Link packets that were unsuccessful in the last test. Failures indicate connectivity or network performance problems (such as overloaded links or devices).

Destination IP/MAC Address is the network address of the target, or destination, device to which you want to test a connection with the switch. An IP address is in the X.X.X.X format where X is a decimal number between 0 and 255. A MAC address is made up of 12 hexadecimal digits, for example, 0060b0-080400.

Number of Packets to Send is the number of times you want the switch to attempt to test a connection.

Timeout in Seconds is the number of seconds to allow per attempt to test a connection before determining that the current attempt has failed.

To halt a Link or Ping test before it concludes, click on the Stop button. **To reset the screen** to its default settings, click on the Defaults button.

CLI: Ping or Link Tests

Ping Tests. You can issue single or multiple ping tests with varying repetitions and timeout periods. The defaults and ranges are:

- Repetitions: 1 (1 - 999)
- Timeout: 5 seconds (1 - 256 seconds)

Syntax: ping <*ip-address*> [repetitions <1 - 999>] [timeout <1 - 256>]

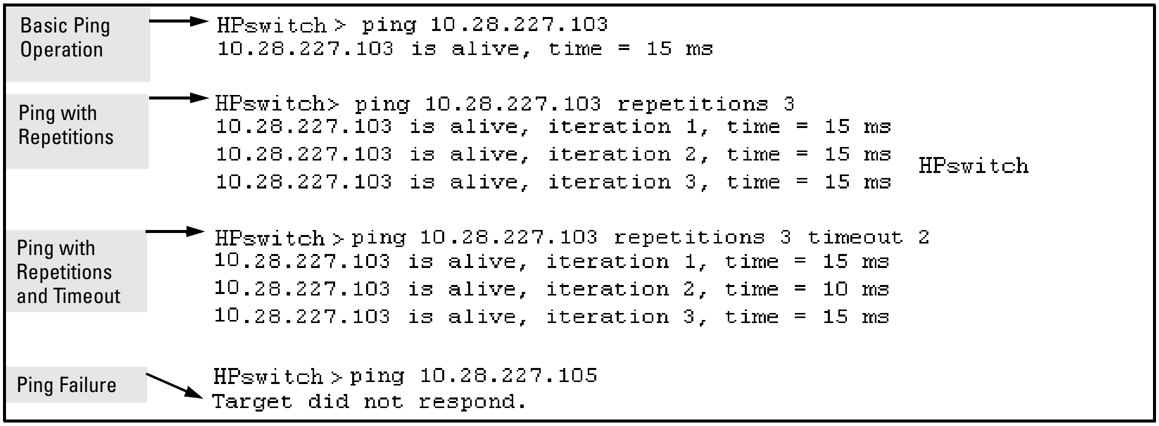

```
HPswitch > ping 10.28.227.103
10.28.227.103 is alive, time = 15 ms

HPswitch> ping 10.28.227.103 repetitions 3
10.28.227.103 is alive, iteration 1, time = 15 ms
10.28.227.103 is alive, iteration 2, time = 15 ms    HPswitch
10.28.227.103 is alive, iteration 3, time = 15 ms

HPswitch >ping 10.28.227.103 repetitions 3 timeout 2
10.28.227.103 is alive, iteration 1, time = 15 ms
10.28.227.103 is alive, iteration 2, time = 10 ms
10.28.227.103 is alive, iteration 3, time = 15 ms

HPswitch >ping 10.28.227.105
Target did not respond.
```

Figure C-7. Examples of Ping Tests

To halt a ping test before it concludes, press Ctrl C.

Link Tests. You can issue single or multiple link tests with varying repetitions and timeout periods. The defaults are:

- Repetitions: 1 (1 - 999)
- Timeout: 5 seconds (1 - 256 seconds)

Syntax: link <*mac-address*> [repetitions <1 - 999>] [timeout <1 - 256>]
[vlan <*vlan-id*>]

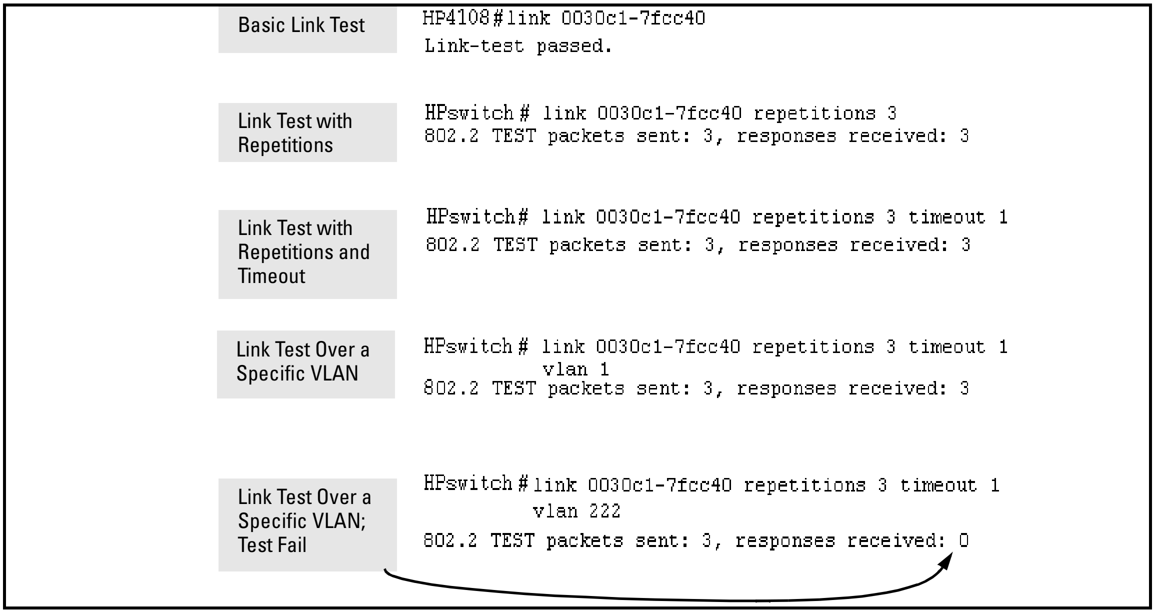

```
HP4108#link 0030c1-7fcc40
Link-test passed.
```

Link Test with Repetitions

```
HPswitch # link 0030c1-7fcc40 repetitions 3
802.2 TEST packets sent: 3, responses received: 3
```

Link Test with Repetitions and Timeout

```
HPswitch# link 0030c1-7fcc40 repetitions 3 timeout 1
802.2 TEST packets sent: 3, responses received: 3
```

Link Test Over a Specific VLAN

```
HPswitch # link 0030c1-7fcc40 repetitions 3 timeout 1
          vlan 1
802.2 TEST packets sent: 3, responses received: 3
```

Link Test Over a Specific VLAN; Test Fail

```
HPswitch #link 0030c1-7fcc40 repetitions 3 timeout 1
          vlan 222
802.2 TEST packets sent: 3, responses received: 0
```

Figure C-8. Example of Link Tests

Displaying the Configuration File

The complete switch configuration is contained in a file that you can browse from either the web browser interface or the CLI. It may be useful in some troubleshooting scenarios to view the switch configuration.

CLI: Viewing the Configuration File

Using the CLI, you can display either the running configuration or the startup configuration. (For more on these topics, see appendix C, "Switch Memory and Configuration".)

Syntax:	write terminal	Displays the running configuration.
	show config	Displays the startup configuration.
	show running-config	Displays the running-config file.

Web: Viewing the Configuration File

To display the running configuration, through the web browser interface:

1. Click on the **Diagnostics** tab.
2. Click on Configuration Report
3. Use the right-side scroll bar to scroll through the configuration listing.

Listing Switch Configuration and Operation Details

The **show tech** command outputs, in a single listing, switch operating and running configuration details from several internal switch sources, including:

- Image stamp (software version data)
- Running configuration
- Event Log listing
- Boot History
- Port settings
- Status and counters — port status
- IP routes
- Status and counters — VLAN information
- GVRP support
- Load balancing (trunk and LACP)

Syntax: show tech

Executing **show tech** outputs a data listing to your terminal emulator. However, using your terminal emulator’s text capture features, you can also save **show tech** data to a text file for viewing, printing, or sending to an associate. For example, if your terminal emulator is the Hyperterminal application available with Microsoft® Windows® software, you can copy the show tech output to a file and then use either Microsoft Word or Notepad to display the data. (In this case, Microsoft Word provides the data in an easier-to-read format.)

To Copy show tech output to a Text File. This example uses the Microsoft Windows terminal emulator. To use another terminal emulator application, refer to the documentation provided with that application.

1. In Hyperterminal, click on **Transfer | Capture Text...**

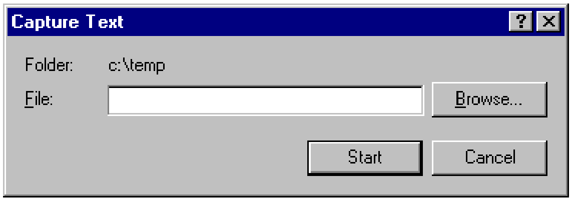

Folder: c:\temp

File:

Figure C-9. The Capture Text window of the Hyperterminal Application

2. In the **File** field, enter the path and file name under which you want to store the **show tech** output.

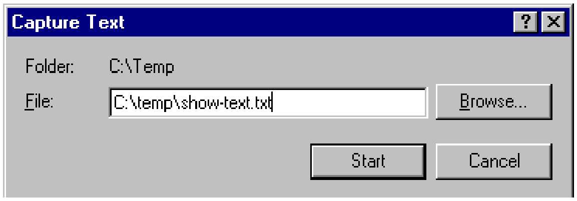

Folder: C:\Temp

File: C:\temp\show-text.txt

Figure C-10. Example of a Path and Filename for Creating a Text File from show tech Output

3. Click [Start] to create and open the text file.
4. Execute **show tech**:

```
HPswitch# show tech
```

   a. Each time the resulting listing halts and displays `-- MORE --`, press the Space bar to resume the listing.

   b. When the CLI prompt appears, the show tech listing is complete. At this point, click on **Transfer | Capture Text | Stop** in HyperTerminal to stop copying data into the text file created in the preceding steps.

Note

Remember to do the above step to stop HyperTerminal from copying into the text file. Otherwise, the text file remains open to receiving additional data from the HyperTerminal screen.

5. To access the file, open it in Microsoft Word, Notepad, or a similar text editor.

CLI Administrative and Troubleshooting Commands

These commands provide information or perform actions that you may find helpful in troubleshooting operating problems with the switch.

Note

For more on the CLI, refer to chapter 3, "Using the Command Line Reference (CLI)".

Syntax:	show version	Shows the software version currently running on the switch, and the flash image from which the switch booted (primary or secondary).
	show boot-history	Displays the switch shutdown history.
	show history	Displays the current command history.
	[no] page	Toggles the paging mode for display commands between continuous listing and per-page listing.
	setup	Displays the Switch Setup screen from the menu interface
	repeat	Repeatedly executes the previous command until a key is pressed.
	kill	Terminates all other active sessions.

Restoring the Factory-Default Configuration

As part of your troubleshooting process, it may become necessary to return the switch configuration to the factory default settings. This process momentarily interrupts the switch operation, clears any passwords, clears the console event log, resets the network counters to zero, performs a complete self test, and reboots the switch into its factory default configuration including deleting an IP address. There are two methods for resetting to the factory-default configuration:

- CLI
- Clear/Reset button combination

Note HP recommends that you save your configuration to a TFTP server before resetting the switch to its factory-default configuration. You can also save your configuration via Xmodem, to a directly connected PC.

CLI: Resetting to the Factory-Default Configuration

This command operates at any level *except* the Operator level.

Syntax: `erase startup-configuration` Deletes the startup-config file in flash so that the switch will reboot with its factory-default configuration.

Note The **erase startup-config** command does not clear passwords.

Clear/Reset: Resetting to the Factory-Default Configuration

To execute the factory default reset, perform these steps:

1. Using pointed objects, simultaneously press both the Reset and Clear buttons on the front of the switch.
2. Continue to press the Clear button while releasing the Reset button.
3. When the Self Test LED begins to flash, release the Clear button.

   The switch will then complete its self test and begin operating with the configuration restored to the factory default settings.

Restoring a Flash Image

The switch can lose its operating system if either the primary or secondary flash image location is empty or contains a corrupted OS file and an operator uses the **erase flash** command to erase a good OS image file from the opposite flash location.

To Recover from an Empty or Corrupted Flash State. Use the switch's console serial port to connect to a workstation or laptop computer that has the following:

- A terminal emulator program with Xmodem capability, such as the HyperTerminal program included in Windows PC software.
- A copy of a good OS image file for the switch.

Note

The following procedure requires the use of Xmodem, and copies an OS image into primary flash only.

This procedure assumes you are using HyperTerminal as your terminal emulator. If you use a different terminal emulator, you may need to adapt this procedure to the operation of your particular emulator.

1. Start the terminal emulator program.
2. Ensure that the terminal program is configured as follows:
   - Baud rate: 9600
   - No parity
   - 8 Bits
   - 1 stop bit
   - No flow control
3. Use the Reset button to reset the switch. The following prompt should then appear in the terminal emulator:

```
Enter h or ? for help.

=>
```

4. Since the OS file is larage, you can increase the speed of the download by changing the switch console and terminal emulator baud rates to a high speed. For example:

   a. Change the switch baud rate to 115,200 Bps.

   ```
   => sp 115200
   ```

   b. Change the terminal emulator baud rate to match the switch speed:

      i. In HyperTerminal, select **Call | Disconnect**.
      ii. Select **File | Properties**.
      iii. click on Configure . . ..
      iv. Change the baud rate to **115200**.
      v. Click on OK. In the next window, click on OK again.
      vi. Select **Call | Connect**
      vii. Press Enter one or more times to display the => prompt.

5. Start the Console Download utility by typing **do** at the => prompt and pressing Enter:

   ```
   => do
   ```

6. You will then see this prompt:

   ```
   You have invoked the console download utility.
   Do you wish to continue? (Y/N)>_
   ```

7. At the above prompt:

   a. Type **y** (for Yes)

   b. Select **Transfer | File** in HyperTerminal.

   c. Enter the appropriate filename and path for the OS image.

   d. Select the **Xmodem** protocol (and not the 1k Xmodem protocol).

   e. Click on Send.

   If you are using HyperTerminal, you will see a screen similar to the following to indicate that the download is in progress:

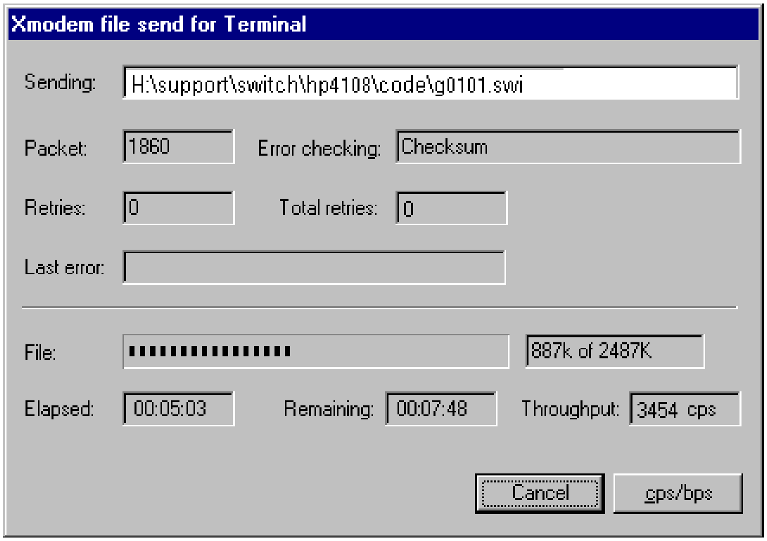

Xmodem file send for Terminal

Sending: H:\support\switch\hp4108\code\g0101.swi

Packet: 1860 Error checking: Checksum

Retries: 0 Total retries: 0

Last error:

File: 887k of 2487K

Elapsed: 00:05:03 Remaining: 00:07:48 Throughput: 3454 cps

Figure C-11. Example of Xmodem Download in Progress

8. When the download completes, the switch reboots from primary flash using the OS image you downloaded in the preceding steps, plus the most recent startup-config file.

MAC Address Management

Overview

The switch assigns MAC addresses in these areas:

- For management functions, one Base MAC address is assigned to the default VLAN (VID = 1). (All VLANs on the switch use the same MAC address.)
- For internal switch operations: One MAC address per port (See "CLI: Viewing the Port and VLAN MAC Addresses" on page D-4.)

MAC addresses are assigned at the factory. The switch automatically implements these addresses for VLANs and ports as they are added to the switch.

Note

The switch's base MAC address is also printed on a label affixed to the back of the switch.

Determining MAC Addresses

MAC Address Viewing Methods

Feature	Default	Menu	CLI	Web
view switch's base (default vlan) MAC address and the addressing for any added VLANs	n/a	D-3	D-4	—
view port MAC addresses (hexadecimal format)	n/a	—	D-4	—

- **Use the menu interface** to view the switch's base MAC address and the MAC address assigned to any non-default VLAN you have configured on the switch.

Note

The switch's base MAC address is used for the default VLAN (VID = 1) that is always available on the switch.

- **Use the CLI** to view the switch's port MAC addresses in hexadecimal format.

Menu: Viewing the Switch's MAC Addresses

The Management Address Information screen lists the MAC addresses for:

- Base switch (default VLAN; VID = 1)
- Any additional VLANs configured on the switch.

Also, the Base MAC address appears on a label on the back of the switch.

Note

The Base MAC address is used by the first (default) VLAN in the switch. This is usually the VLAN named "DEFAULT_VLAN" unless the name has been changed (by using the VLAN Names screen). On the Series 5300XL switch, the VID (VLAN identification number) for the default VLAN is always "1", *and cannot be changed*.

To View the MAC Address (and IP Address) assignments for VLANs Configured on the Switch:

1. From the Main Menu, Select

   1. Status and Counters

   2. Switch Management Address Information

   If the switch has only the default VLAN, the following screen appears. If the switch has multiple static VLANs, each is listed with its address data.

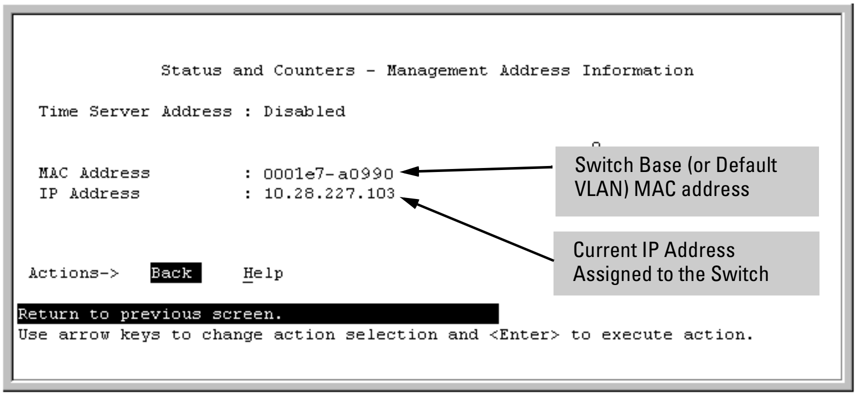

```
                Status and Counters - Management Address Information

  Time Server Address : Disabled


  MAC Address         : 0001e7-a0990
  IP Address          : 10.28.227.103


 Actions->   Back    Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Switch Base (or Default VLAN) MAC address

Current IP Address Assigned to the Switch

Figure B-1. Example of the Management Address Information Screen

CLI: Viewing the Port and VLAN MAC Addresses

The MAC address assigned to each switch port is used internally by such features as Flow Control and the Spanning Tree Protocol. Using the **walkmib** command to determine the MAC address assignments for individual ports can sometimes be useful when diagnosing switch operation.

The switch allots 26 MAC addresses per slot. For a given slot, if a four-port module is installed, then the switch uses the first four MAC addresses in the allotment for that slot, and the remaining 22 MAC addresses are unused. If a 24-port module is installed, the switch uses the first 24 MAC addresses in the allotment, and so-on. The switch's base MAC address is assigned to VLAN (VID) 1 and appears in the **walkmib** listing after the MAC addresses for the ports. (All VLANs in the switch have the same MAC address.)

To display the switch's MAC addresses, use the **walkmib** command at the command prompt:

Note

This procedure displays the MAC addresses for all ports and existing VLANs in the switch, regardless of which VLAN you select.

1. If the switch is at the CLI Operator level, use the **enable** command to enter the Manager level of the CLI.
2. Type the following command to display the MAC address for each port on the switch:

   ```
   HPswitch# walkmib ifPhysAddress
   ```

   (The above command is not case-sensitive.)

For example, with a 4-port module in slot A and a 24-port module in slot B, and three VLANs present:

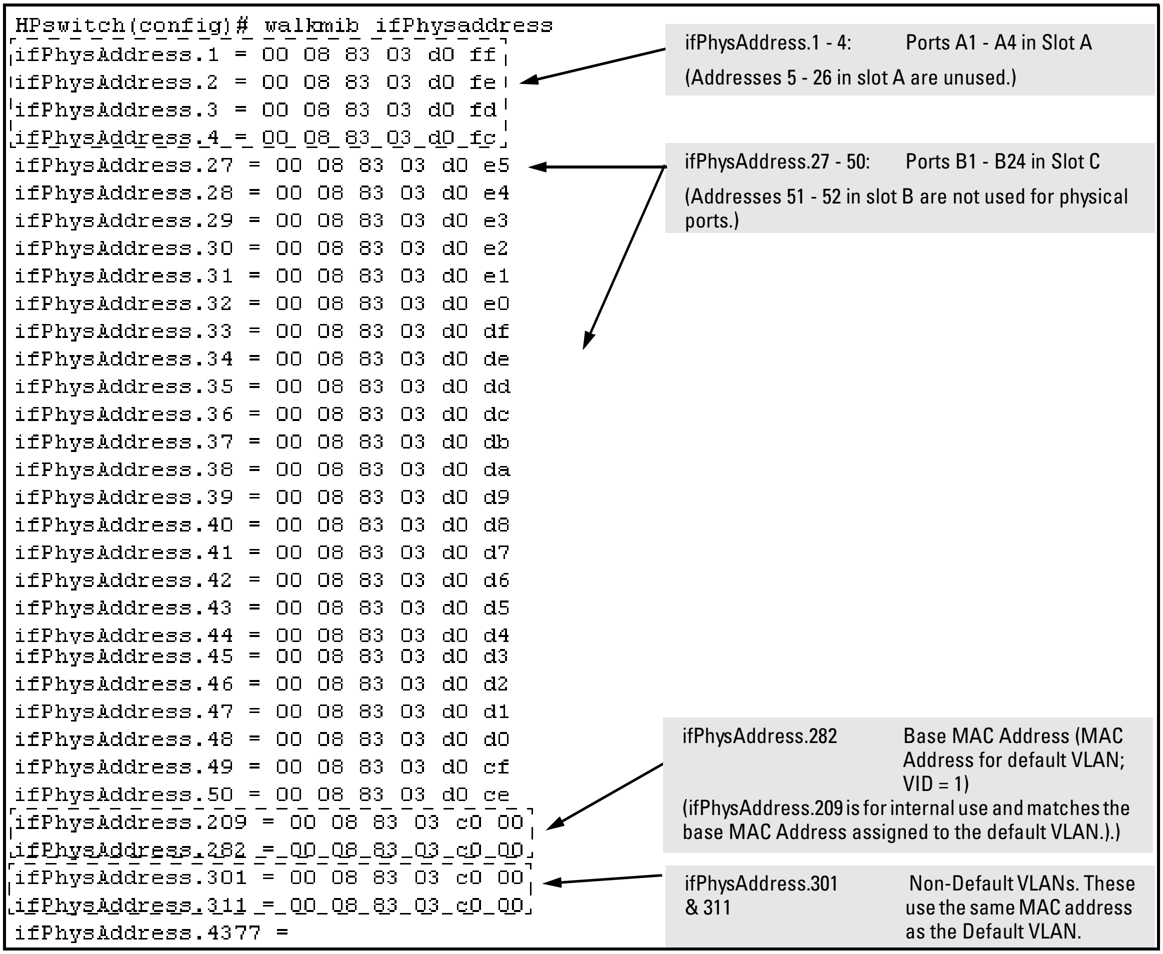

```
HPswitch(config)# walkmib ifPhysaddress
ifPhysAddress.1 = 00 08 83 03 d0 ff
ifPhysAddress.2 = 00 08 83 03 d0 fe
ifPhysAddress.3 = 00 08 83 03 d0 fd
ifPhysAddress.4 = 00 08 83 03 d0 fc
ifPhysAddress.27 = 00 08 83 03 d0 e5
ifPhysAddress.28 = 00 08 83 03 d0 e4
ifPhysAddress.29 = 00 08 83 03 d0 e3
ifPhysAddress.30 = 00 08 83 03 d0 e2
ifPhysAddress.31 = 00 08 83 03 d0 e1
ifPhysAddress.32 = 00 08 83 03 d0 e0
ifPhysAddress.33 = 00 08 83 03 d0 df
ifPhysAddress.34 = 00 08 83 03 d0 de
ifPhysAddress.35 = 00 08 83 03 d0 dd
ifPhysAddress.36 = 00 08 83 03 d0 dc
ifPhysAddress.37 = 00 08 83 03 d0 db
ifPhysAddress.38 = 00 08 83 03 d0 da
ifPhysAddress.39 = 00 08 83 03 d0 d9
ifPhysAddress.40 = 00 08 83 03 d0 d8
ifPhysAddress.41 = 00 08 83 03 d0 d7
ifPhysAddress.42 = 00 08 83 03 d0 d6
ifPhysAddress.43 = 00 08 83 03 d0 d5
ifPhysAddress.44 = 00 08 83 03 d0 d4
ifPhysAddress.45 = 00 08 83 03 d0 d3
ifPhysAddress.46 = 00 08 83 03 d0 d2
ifPhysAddress.47 = 00 08 83 03 d0 d1
ifPhysAddress.48 = 00 08 83 03 d0 d0
ifPhysAddress.49 = 00 08 83 03 d0 cf
ifPhysAddress.50 = 00 08 83 03 d0 ce
ifPhysAddress.209 = 00 08 83 03 c0 00
ifPhysAddress.282 = 00 08 83 03 c0 00
ifPhysAddress.301 = 00 08 83 03 c0 00
ifPhysAddress.311 = 00 08 83 03 c0 00
ifPhysAddress.4377 =
```

ifPhysAddress.1 - 4: Ports A1 - A4 in Slot A

(Addresses 5 - 26 in slot A are unused.)

ifPhysAddress.27 - 50: Ports B1 - B24 in Slot C

(Addresses 51 - 52 in slot B are not used for physical ports.)

ifPhysAddress.282 Base MAC Address (MAC Address for default VLAN; VID = 1)
(ifPhysAddress.209 is for internal use and matches the base MAC Address assigned to the default VLAN.).)

ifPhysAddress.301 & 311 Non-Default VLANs. These use the same MAC address as the Default VLAN.

Figure B-2. Example of Port MAC Address Assignments

Daylight Savings Time on HP Procurve Switches

This information applies to the following HP Procurve switches:

- 2512
- 2524
- 4104GL
- 4108GL
- 5304XL
- 5308XL
- 1600M
- 2400M
- 2424M
- 4000M
- 8000M
- 212M
- 224M
- HP AdvanceStack Switches
- HP AdvanceStack Routers

HP Procurve switches provide a way to automatically adjust the system clock for Daylight Savings Time (DST) changes. To use this feature you define the month and date to begin and to end the change from standard time. In addition to the value “none” (no time changes), there are five pre-defined settings, named:

- Alaska
- Canada and Continental US
- Middle Europe and Portugal
- Southern Hemisphere
- Western Europe

The pre-defined settings follow these rules:

Alaska:

- Begin DST at 2am the first Sunday on or after April 24th.
- End DST at 2am the first Sunday on or after October 25th.

Canada and Continental US:

- Begin DST at 2am the first Sunday on or after April 1st.
- End DST at 2am the first Sunday on or after October 25th.

Middle Europe and Portugal:

- Begin DST at 2am the first Sunday on or after March 25th.
- End DST at 2am the first Sunday on or after September 24th.

Southern Hemisphere:

- Begin DST at 2am the first Sunday on or after October 25th.
- End DST at 2am the first Sunday on or after March 1st.

Western Europe:

- Begin DST at 2am the first Sunday on or after March 23rd.
- End DST at 2am the first Sunday on or after October 23rd.

A sixth option named "User defined" allows you to customize the DST configuration by entering the beginning month and date plus the ending month and date for the time change. The menu interface screen looks like this (all month/ date entries are at their default values):

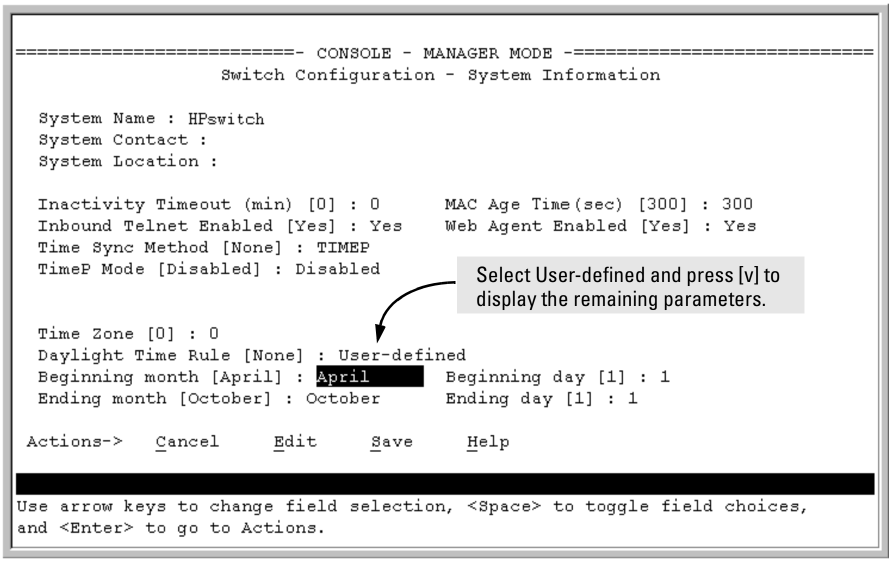

```
==========================- CONSOLE - MANAGER MODE -============================
                   Switch Configuration - System Information

  System Name : HPswitch
  System Contact :
  System Location :

  Inactivity Timeout (min) [0] : 0      MAC Age Time(sec) [300] : 300
  Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes
  Time Sync Method [None] : TIMEP
  TimeP Mode [Disabled] : Disabled


  Time Zone [0] : 0
  Daylight Time Rule [None] : User-defined
  Beginning month [April] : April       Beginning day [1] : 1
  Ending month [October] : October      Ending day [1] : 1

 Actions->   Cancel     Edit     Save     Help


Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Select User-defined and press [v] to display the remaining parameters.

Figure E-1. Menu Interface with "User-Defined" Daylight Time Rule Option

Before configuring a “User defined” Daylight Time Rule, it is important to understand how the switch treats the entries. The switch knows which dates are Sundays, and uses an algorithm to determine on which date to change the system clock, given the configured “Beginning day” and “Ending day”:

- If the configured day is a Sunday, the time changes at 2am on that day.
- If the configured day is not a Sunday, the time changes at 2am on the first Sunday after the configured day.

This is true for both the “Beginning day” and the “Ending day”.

With that algorithm, one should use the value “1” to represent “first Sunday of the month”, and a value equal to “number of days in the month minus 6" to represent “last Sunday of the month”. This allows a single configuration for every year, no matter what date is the appropriate Sunday to change the clock.

Index

invent

Technical information in this document
is subject to change without notice.

Edition 4, October 2002

Manual Part Number
5990-3016